

Not All Tokens Matter: Towards Efficient LLM Reasoning via Token Significance in Reinforcement Learning

Hanbing Liu^{1*†} Lang Cao^{2*} Yuanyi Ren^{3*} Mengyu Zhou^{4‡} Haoyu Dong⁵
Xiaojun Ma⁵ Shi Han⁵ Dongmei Zhang⁵

¹Tsinghua University ²University of Illinois Urbana-Champaign ³Peking University
⁴Qwen Large Model Application Team, Alibaba ⁵Microsoft
liuhb24@mails.tsinghua.edu.cn, langcao2@illinois.edu, yyren@pku.edu.cn

Abstract

Large language models (LLMs) show strong reasoning abilities but often produce unnecessarily long explanations that reduce efficiency. Although reinforcement learning (RL) has been used to improve reasoning, most methods focus on accuracy and rely on uniform length-based rewards that overlook the differing contributions of individual tokens, often harming correctness. We revisit length optimization in RL through the perspective of *token significance*. Observing that many chain-of-thought (CoT) tokens contribute little to the final answer, we introduce a significance-aware length reward that selectively penalizes insignificance tokens, reducing redundancy while preserving essential reasoning. We also propose a dynamic length reward that encourages more detailed reasoning early in training and gradually shifts toward conciseness as learning progresses. Integrating these components into standard policy optimization yields a framework that improves both reasoning efficiency and accuracy. Experiments across multiple benchmarks demonstrate substantial reductions in response length while preserving or improving correctness, highlighting the importance of modeling token significance for efficient LLM reasoning. Code is available at <https://github.com/microsoft/Bingo>.

1 Introduction

Large language models (LLMs) (OpenAI, 2024; Gunasekar et al., 2023) have demonstrated impressive reasoning capabilities across a variety of tasks, from arithmetic problem solving (Uesato et al., 2022; Hendrycks et al., 2021; Veeraboina, 2023)

*Work done during internship at Microsoft.

†Affiliated with the Shenzhen Key Laboratory of Ubiquitous Data Enabling, Tsinghua Shenzhen International Graduate School, Tsinghua University.

‡Corresponding author. zhoumengyu.zmy@alibaba-inc.com.

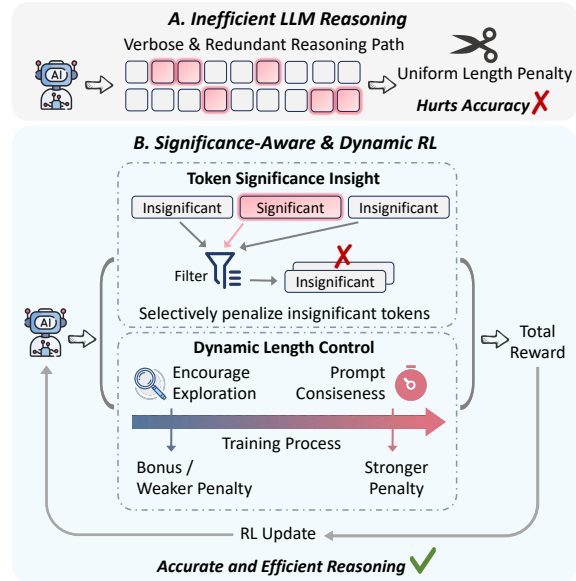


Figure 1: (A) **Limitations of Uniform Length Penalization.** LLMs often produce verbose and redundant reasoning, and applying a uniform length penalty, as done in many prior RL approaches, fails to account for the differing importance of individual tokens, which can lead to accuracy degradation; (B) **Reinforcement Learning with Token Significance and Dynamic Length Control.** Our method models token significance to selectively penalize unimportant tokens and introduces dynamic length control to balance exploration and conciseness throughout training, enabling LLMs to generate reasoning that is both accurate and efficient.

to commonsense reasoning (Chen et al., 2023). A key observation from recent work is that sufficiently large models can exhibit emergent reasoning abilities, such as chain-of-thought (CoT) reasoning (Wei et al., 2022b), without explicit supervision (Wei et al., 2022a; Suzgun et al., 2022). Despite these successes, a major challenge persists: LLMs often generate unnecessarily verbose or redundant reasoning traces, leading to inefficiencies in computational cost, redundancy, and latency.

Improving reasoning efficiency of LLMs has thus emerged as an important research direc-

tion (Qu et al., 2025; Sui et al., 2025; Li et al., 2025; Wang et al., 2025). Prior work in this area can be broadly categorized into supervised fine-tuning (SFT) approaches (Xia et al., 2025; Xu et al., 2025a; Zhang et al., 2025; Kang et al., 2024) and reinforcement learning (RL) approaches (Luo et al., 2025; Team et al., 2025; Arora and Zanette, 2025; Aggarwal and Welleck, 2025; Yeo et al., 2025; Shi et al., 2025). SFT-based methods focus on constructing compressed reasoning traces and training models to imitate them. While these approaches can be effective, they rely on high-quality compressed supervision, which is costly to obtain and often lacks generalizability across diverse tasks. RL-based methods typically introduce length-based rewards that penalize overly long responses to encourage brevity. However, the design of such rewards or penalties in RL-based methods remains underexplored and is often overly simplistic. For example, O1-Pruner (Luo et al., 2025) applies a uniform penalty to all samples, assuming that every response should be shortened. This assumption often leads to performance degradation, as not all reasoning traces are equally verbose—some require more detailed steps to arrive at the correct answer. To address this, other works have proposed more selective penalty strategies, conditioning penalties on sample correctness (Qu et al., 2025; Team et al., 2025; Yeo et al., 2025) or estimated difficulty (Shi et al., 2025). These approaches typically assign stronger penalties to simpler questions and weaker ones to more challenging cases. However, accurately estimating question difficulty remains a fundamental challenge, and unresolved hard questions often lead to unnecessarily long responses, further undermining reasoning efficiency.

Despite growing interest, current designs of length-based rewards remain limited, as they often fail to adequately promote concise reasoning while preserving answer accuracy. For example, prior work has largely overlooked the impact of token-level contributions on the overall efficiency of reasoning as illustrated in Figure 1 (A). In this work, we approach the problem from a novel perspective grounded in the concept of *token significance*. Our motivation arises from observed token redundancy in LLMs (Hou et al., 2022; Lin et al., 2025), where many tokens in chain-of-thought (CoT) reasoning contribute little to the final answer. We posit that *not all tokens are equally important for efficient reasoning*—many are insignificant, such as redundant phrases or unnecessary intermediate steps, and

can be removed without degrading performance. Existing reward designs often overlook this distinction. In contrast, we introduce a *significance-aware length reward* that selectively penalizes only those insignificant tokens which do not meaningfully contribute to the final answer, while preserving essential reasoning steps.

We also observe that effectively handling hard questions is essential for efficient reasoning. Prior work (Muennighoff et al., 2025; Wu et al., 2025) has shown that encouraging extended CoT reasoning can improve performance by enabling deeper exploration, which may help solve more difficult questions. Therefore, it is intuitive to use length as an incentive for hard questions. However, LLMs should solve difficult questions not only accurately but also concisely. Applying a static length incentive can lead to unnecessarily long responses, which may still fail to produce correct answers. To address this, we incorporate a *dynamic length reward* that adapts over the course of training. This reward is applied to significant tokens in incorrect samples to balance exploration and efficiency. Specifically, it encourages longer reasoning in the early training phase to promote exploration, and gradually shifts toward penalizing excessive length in later stages to promote conciseness.

Building on these insights, we introduce **BINGO** (**B**oosting **E**fficient Reason**I**NG in Policy **O**ptimization), an RL framework that incorporates our two proposed reward mechanisms into standard RL algorithms such as Proximal Policy Optimization (PPO) (Schulman et al., 2017) as illustrated in Figure 1 (B). This enables joint optimization of both reasoning accuracy and efficiency. Extensive experiments across diverse reasoning benchmarks show that BINGO consistently outperforms strong baselines by reducing redundant computation while maintaining or improving accuracy (Appendix I).

In summary, this paper makes the following key contributions:

- **Token Significance Insight.** We introduce the concept of *token significance* in policy optimization, distinguishing between *significant* and *insignificant* tokens in reasoning traces. This insight motivates our *significance-aware length reward*, which explicitly penalizes uninformative tokens while preserving critical reasoning content, enabling more targeted and effective length control.
- **Dynamic Length Control.** We propose a

dynamic length reward strategy that adjusts the reward signal over the course of training—encouraging longer reasoning in the early stages to foster exploration, and gradually promoting conciseness as the model converges.

- **Efficiency-Oriented RL Framework.** We develop **BINGO**, a new reinforcement learning framework that integrates both reward strategies. Extensive experiments across multiple reasoning benchmarks, along with comprehensive analyses, demonstrate its effectiveness.

2 Related Work

Reinforcement Learning for Large Language Models. Reinforcement Learning (RL) (Kaelbling et al., 1996) has emerged as a powerful paradigm for aligning large language models (LLMs) with human preferences. In Reinforcement Learning from Human Feedback (RLHF) (Christiano et al., 2017; Stiennon et al., 2020; Ouyang et al., 2022), the Proximal Policy Optimization (PPO) algorithm (Schulman et al., 2017) is employed alongside human preference data to train a reward model that steers the fine-tuning of LLMs. Building on PPO, subsequent works like GRPO (Shao et al., 2024) and REINFORCE++ (Hu, 2025) have proposed improved variants to address its limitations. Beyond alignment, RL has also shown promise in improving the reasoning capabilities of LLMs. Early studies (Lightman et al., 2023; Uesato et al., 2022) demonstrated that reward-guided training can enhance multi-step reasoning performance. More recently, DeepSeek-R1 (DeepSeek-AI et al., 2025) demonstrated that large-scale RL can substantially boost reasoning ability across a wide range of tasks, pointing to a promising direction for future work.

Efficient Reasoning with Large Language Models. Recent advances have empowered language models to perform strong reasoning via inference-time techniques such as chain-of-thought prompting (Wei et al., 2023; Yao et al., 2023; Cao, 2024; Wang et al., 2023) and post-training (Lightman et al., 2023; Uesato et al., 2022; DeepSeek-AI et al., 2025; Jiang et al., 2025; Cao et al., 2025). More recent work has shifted to optimizing both accuracy and efficiency. Some approaches improve efficiency at inference time, such as token-budget-aware reasoning (Han et al., 2025), or prompting strategies like “reason-without-thinking”

(Ma et al., 2025) and chain-of-draft (Xu et al., 2025b). Others apply post-training optimization via supervised fine-tuning (SFT), including TokenSkip (Xia et al., 2025), TwT (Xu et al., 2025a), LightThinker (Zhang et al., 2025), and C3oT (Kang et al., 2024). These SFT methods primarily construct high-quality compressed reasoning paths containing key information, and train the models on them. In parallel, RL-based approaches often improve efficiency by incorporating length controls or penalties into their reward functions (Luo et al., 2025; Team et al., 2025; Arora and Zanette, 2025; Aggarwal and Welleck, 2025; Yeo et al., 2025; Shi et al., 2025). For instance, O1-Pruner (Luo et al., 2025) uses offline length rewards comparing samples against mean lengths. Kimi k1.5 (Team et al., 2025) applies online penalties to correct samples only. Building on prior RL-based approaches, we advance length-based reward design to enable LLMs to balance reasoning accuracy with computational efficiency.

3 Methodology

In this section, we introduce the design of the significance-aware length reward and the dynamic length reward, and explain how these two reward mechanisms are integrated into the BINGO framework, as illustrated in Figure 2. All notations are listed at Appendix R.

3.1 Task Formulation

Chain of Thought Reasoning. Let x denote a prompt, and let $y = (y_1, y_2, \dots, y_n)$ represent the sequence generated by a language model parameterized by θ , where y_i is the i -th token in the sequence, and n is the total length of the sequence. Tokens are generated autoregressively from the conditional distribution:

$$\pi_{\theta}(y | x) = \prod_{t=1}^n \pi_{\theta}(y_t | x, y_{1:t-1}), \quad (1)$$

where the product runs over all tokens in the sequence, with each token y_i (i.e., the action a_i) conditioned on the prompt x and the previous tokens $y_{1:i-1}$ (i.e., the state s_i). Generation continues until an end-of-sequence (EOS) token is produced, signaling the completion of the response. During this process, the model may produce intermediate reasoning tokens, referred to as a *chain of thought* (CoT) (Wei et al., 2022b), before generating the final answer. Therefore, the full output sequence,

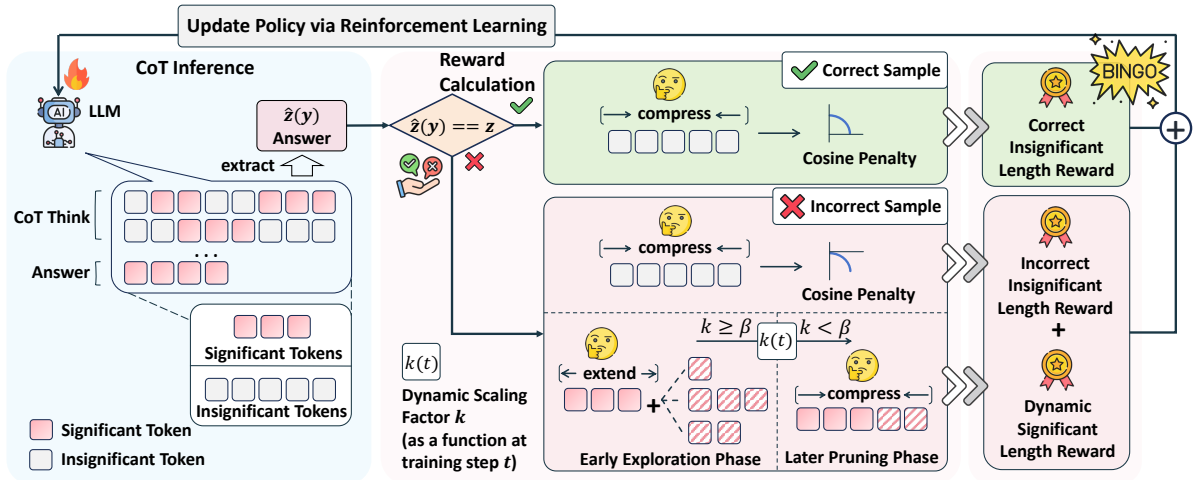


Figure 2: **Illustration of the BINGO framework.** Given a generated CoT trace, the LLM first distinguishes between *significant* and *insignificant* tokens. A dynamic length reward is then computed based on token type and sample correctness. During the early exploration phase of training ($k(t) \geq \beta$), the reward encourages extended reasoning for significant tokens in incorrect samples while penalizing insignificant tokens in all cases. As training progresses ($k(t) < \beta$), the reward shifts toward promoting conciseness by discouraging both significant and insignificant length where appropriate. This two-stage strategy allows the model to first explore broadly and then compress effectively. The aggregated rewards are then used to update the policy via RL, resulting in more accurate and efficient reasoning.

denoted as y , consists of both the chain of thought and the final answer.

Optimization Objective of Efficient Reasoning.

The performance of the model in efficient reasoning is assessed along two key dimensions: *accuracy* and *efficiency*.

Accuracy is measured by the *Exact Match* (EM) metric, which evaluates whether the model’s final answer matches the ground truth. Let $\hat{z}(y)$ denote the final answer extracted from the model-generated sequence y , typically corresponding to its final segment. Let z be the ground-truth answer. Then EM is defined as:

$$\text{EM} = \mathbb{E}_{x \sim \phi} \mathbb{E}_{y \sim \pi_{\theta}(\cdot|x)} \mathbb{1} [\hat{z}(y) = z], \quad (2)$$

where ϕ denotes the distribution over prompts. The indicator function returns 1 if the predicted answer exactly matches the ground truth, and 0 otherwise.

Efficiency is measured by the *response length* L , typically defined as the number of tokens n in the generated sequence $y = (y_1, y_2, \dots, y_n)$. While longer sequences may offer detailed reasoning, they often result in higher computational cost. Thus, reducing unnecessary tokens without harming accuracy is crucial for practical deployment. An ideal model achieves high EM while minimizing the average response length L , striking a balance between correctness and conciseness.

Token Significance. CoT responses consist of a sequence of tokens that contribute unequally to the

overall reasoning process. This observation motivates a token-level view of reasoning efficiency, in which the contribution of each token is described by the notion of *token significance*.

Given a generated sequence $y = (y_1, y_2, \dots, y_n)$, the significance of a token y_i is defined via a contextual scoring function:

$$S(y_i) = F(y_i | \mathbf{y}_{\leq n}), \quad (3)$$

where $F(\cdot)$ assigns a scalar score to a token based on its surrounding context within the full sequence, and $\mathbf{y}_{\leq n}$ denotes the complete set of tokens in the generated response. Tokens associated with higher scores are considered to have higher significance. Details of the scoring function and a discussion of alternative definitions of token significance are presented in Appendix C.

3.2 Significance-Aware Length Reward

To enhance the efficiency of CoT generation, it is crucial to recognize that not all tokens in a CoT sequence contribute equally to deriving the final answer. Tokens with low importance scores are considered insignificant, while those with high scores are deemed significant. Specifically, we classify tokens as follows:

$$\text{Token } y_i \text{ is } \begin{cases} \text{insignificant,} & \text{if } S(y_i) < \tau, \\ \text{significant,} & \text{if } S(y_i) \geq \tau. \end{cases} \quad (4)$$

We then compute the total number of significant tokens L^s and insignificant tokens L^{is} in the response

as:

$$L^s = \sum_{i=1}^n \mathbb{1}(S(y_i) \geq \tau), L^{is} = \sum_{i=1}^n \mathbb{1}(S(y_i) < \tau), \quad (5)$$

where $\mathbb{1}(\cdot)$ is the indicator function, and τ is a pre-defined threshold. To encourage brevity while maintaining reasoning quality, we introduce a **significance-aware length reward** that penalizes the excessive use of insignificant tokens through a cosine-based decay:

$$r_{is}(y) = \cos\left(\text{clip}\left(\frac{L^{is}}{L_{\text{ref}}^{is}}, 0, \frac{\pi}{2}\right)\right) + \mathbb{1}[\hat{z}(y) = z] \quad (6)$$

where L_{ref}^{is} denotes the number of insignificant tokens in a reference response. The cosine function ensures a smooth, non-linear penalty that gradually decreases the reward as L^{is} increases, while the clipping operation bounds the angle to the interval $[0, \frac{\pi}{2}]$, preventing negative rewards. The final reward combines this length-based penalty with an answer reward derived from the EM indicator, ensuring that answer correctness is preserved.

This reward formulation ensures that shorter or equally concise responses—measured in terms of insignificant content—receive higher rewards, while excessively verbose outputs are gently penalized. Notably, our approach preserves natural fluency and coherence in generated text by constraining only the aggregate length of insignificant tokens, without dictating specific token selections or sequences in RL-based training. Compared to standard length-based penalties, our significance-aware approach achieves equal or greater length reductions with less accuracy degradation by selectively penalizing insignificant tokens, as theoretically justified in Appendix D.

3.3 Dynamic Length Reward for Significant Tokens

While insignificant tokens are consistently penalized to reduce redundancy, significant tokens warrant a more nuanced approach. In the early stages of training, allowing longer reasoning with significant content can facilitate exploration and support the development of robust problem-solving strategies. However, as training progresses, conciseness becomes increasingly important for improving efficiency.

To accommodate this shift, we introduce a **dynamic length reward** for significant tokens that

evolves over time based on the model’s learning trajectory. This adaptive mechanism is guided by a dynamic scaling factor that captures trends in accuracy and modulates the reward accordingly. Formally, the length-based reward for significant tokens is defined as:

$$r_s(y) = \begin{cases} k \cdot \frac{L^s}{L_{\text{ref}}^s}, & \text{if } k \geq \beta \\ -\alpha \cdot t \cdot \frac{L^s}{L_{\text{ref}}^s}, & \text{if } k < \beta \end{cases} \quad (7)$$

where L^s represents the number of significant tokens in the generated output, L_{ref}^s is the corresponding value from the reference model, and k is a dynamic scaling factor that reflects the reasoning trend during training. The training step t begins at 1 and increments gradually when k first falls below the threshold β , which determines when the model transitions from incentivizing longer significant token lengths to penalizing them. α is a weight that determines the rate of decay in this process. The value of k is estimated by fitting a linear model to recent training steps:

$$k = \frac{\sum_{t=S_a}^{S_b} (t - \bar{t})(acc_t - \overline{acc})}{\sum_{t=S_a}^{S_b} (t - \bar{t})^2} \quad (8)$$

where acc_t denotes the training batch accuracy at training step t , \bar{t} and \overline{acc} are the mean step index and mean accuracy over the interval $[S_a, S_b]$. A positive k indicates an upward accuracy trend, suggesting that the model is still in an improvement phase. As training progresses and accuracy plateaus, k approaches zero or becomes negative. The theoretical rationale behind the design of our dynamic length reward schedule is discussed in detail in Appendix E. This dynamic adaptation allows the model to balance early-stage exploration with late-stage compression, fostering reasoning strategies that are both effective and efficient.

3.4 Boosting Efficient Reasoning in Policy Optimization

We propose a novel reinforcement learning algorithm, BINGO (**B**oosting **E**fficient ReasonING in Policy **O**ptimization), designed to jointly optimize reasoning performance and efficiency. BINGO extends the reinforcement learning framework, primarily based on Proximal Policy Optimization (PPO) in this work, by introducing two key innovations: a **significance-aware length reward** and a **dynamic length reward**.

As discussed in Section 3.2, we begin by categorizing tokens into *significant* and *insignificant*

based on their significance scores. To promote concise yet informative responses, we introduce a cosine-based reward function that adjusts penalties according to the length composition of the response. For correctly answered samples, the reward penalizes only the length of the insignificant portion, reducing verbosity while preserving essential reasoning. For incorrect samples, the reward both penalizes the use of insignificant tokens and encourages the generation of more significant reasoning content.

To balance exploration and efficiency over the course of training, we incorporate a time-dependent mechanism that gradually reduces the incentive for longer responses. As detailed in Section 3.3, this dynamic reward decays as the model converges, shifting the focus from exploration to conciseness.

The overall reward $R^{\text{BINGO}}(y)$ formulation integrates these components into a unified objective:

$$R^{\text{BINGO}}(y) = \begin{cases} \lambda_c \cdot r_{is}(y), & \text{if correct,} \\ \underbrace{\lambda_w^{is} \cdot [r_{is}(y) - 1]}_{\text{Incorrect insignificant length}} + \underbrace{\min(0, r_s(y) - \lambda_w^s)}_{\text{Dynamic significant length}}, & \text{if incorrect.} \end{cases} \quad (9)$$

where the coefficient λ_c controls the strength of the penalty applied to correct responses, while λ_w^{is} determines the magnitude of the penalty for incorrect ones. The parameter λ_w^s serves as a dynamic threshold to balance exploration when the model generates incorrect outputs.

We optimize the policy using the proximal policy optimization objective with the reward R_{BINGO} defined by Equation 9. The surrogate objective is:

$$\mathcal{J}_{\text{BINGO}}(\theta) = \mathbb{E}_t \left[\min \left(r_t(\theta) \hat{A}_t, \text{clip} \left(r_t(\theta), 1 - \epsilon, 1 + \epsilon \right) \hat{A}_t \right) \right], \quad (10)$$

where:

- $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$ is the importance sampling ratio,
- \hat{A}_t is the advantage estimate at time step t , computed via generalized advantage estimation using the final sequence-level reward R^{BINGO} and the value predictions $V(s_t)$.
- ϵ is a clipping parameter.

Therefore, BINGO achieves a favorable trade-off between accuracy and efficiency by maximizing the objective function $\mathcal{J}_{\text{BINGO}}(\theta)$.

4 Experiments

4.1 Experimental Setup and Evaluation Metrics

We fine-tune two reasoning models, *DeepSeek-R1-Distill-Qwen-1.5B* and *DeepSeek-R1-Distill-Qwen-7B*, together with an instruction-tuned model, *Qwen2.5-Math-7B-Instruct*, on the MATH training split. Evaluation is performed on MATH500, GSM8K, AIME2024, and THEOREMQA. MATH500 serves as the in-distribution benchmark, while the remaining datasets constitute out-of-distribution evaluations, covering both mathematical reasoning, and commonsense question answering.

We compare BINGO against a frozen base model, a Vanilla PPO baseline, and several recent state-of-the-art methods. To enable fair comparison, all baselines are re-implemented within a unified PPO framework, isolating the effect of different length-based reward designs. Detailed experimental settings and baseline descriptions are provided in Appendix G. A thorough hyperparameter study can be found in Appendix Q.

We report two model variants: BINGO-A, selected for peak validation accuracy, and BINGO-E, selected for stable response length. This dual-reporting strategy enables practitioners to choose a model variant based on their preference for accuracy or efficiency. They may correspond to the same checkpoint. To evaluate reasoning efficiency, we report accuracy (Acc), response length (Len), and length-normalized accuracy (L-Acc), which jointly measures correctness and conciseness:

$$\text{L-Acc} = \text{Acc} \times \sqrt{1 - \frac{L}{L_{\text{max}}}}. \quad (11)$$

where L is the average response length and L_{max} is the maximum allowable length. A detailed definition and theoretical analysis of L-Acc are provided in Appendix F.

4.2 Performance Comparison with Baseline Methods

BINGO outperforms existing methods in L-Acc:

As shown in Table 1, both BINGO-A and BINGO-E achieve the highest L-Acc across all four benchmarks, outperforming previous baselines such as Vanilla PPO, Efficient Reasoning, and DAST.

BINGO-A improves accuracy while significantly reducing response length: BINGO-A reduces average response length by up to 68% compared

Table 1: **Comparison of different length-based rewards on reasoning benchmarks.** Each method is evaluated using DeepSeek-R1-Distill-Qwen-1.5B as the base model by answer accuracy (Acc, %), response length (Len), and length-normalized accuracy (L-Acc, %). The best performance is highlighted in **dark blue**, and the second-best in **light blue**.

Length-based Reward	MATH500			GSM8K			TheoremQA			AIME2024		
	Acc↑	Len↓	L-Acc↑	Acc↑	Len↓	L-Acc↑	Acc↑	Len↓	L-Acc↑	Acc↑	Len↓	L-Acc↑
Vanilla PPO (Schulman et al., 2017)	81.4	2,771	66.2	85.4	1,310	78.2	32.3	4,146	22.7	26.7	6,961	10.3
O1-Pruner (Luo et al., 2025)	74.4	991	69.8	81.4	211	80.3	28.6	524	27.6	26.7	5,958	13.9
kimi-k1.5 (Team et al., 2025)	80.4	1,692	71.6	85.4	739	81.5	34.4	2,136	29.6	33.3	5,159	20.3
Effi. Reasoning (Arora and Zanette, 2025)	82.6	2,395	69.5	86.4	1,155	80.0	34.8	3,560	26.2	36.7	5,771	19.9
Demystifying (Yeo et al., 2025)	80.2	1,411	73.0	86.6	548	83.6	35.1	1,976	30.6	30.0	6,183	14.9
DAST (Shi et al., 2025)	81.2	1,770	71.9	82.0	456	79.6	35.2	2,325	29.8	36.7	5,400	21.4
<i>Bingo (Ours)</i>												
Bingo-A	82.2	894	77.6	87.0	563	83.9	36.8	1,648	32.9	33.3	2,943	26.7
Bingo-E	80.6	779	76.7	86.7	345	84.9	36.7	1,584	33.0	33.3	2,943	26.7

Table 2: **Performance comparison across model scales and types.** Accuracy (Acc), average output length (Length), and length-normalized accuracy (L-Acc, %) on four benchmarks. The best performance is highlighted in **dark blue**, and the second-best in **light blue**.

Method	MATH500			GSM8K			TheoremQA			AIME2024		
	Acc↑	Len↓	L-Acc↑	Acc↑	Len↓	L-Acc↑	Acc↑	Len↓	L-Acc↑	Acc↑	Len↓	L-Acc↑
<i>DeepSeek-R1-Distill-Qwen-1.5B</i>												
Base	63.2	3,913	45.7	73.2	2,025	63.5	18.7	5,741	10.3	16.7	7,027	6.3
PPO	81.4	2,771	66.2	85.4	1,310	78.2	32.3	4,146	22.7	26.7	6,961	10.3
Bingo-A (Ours)	82.2	894	77.6	87.0	563	83.9	36.8	1,648	32.9	33.3	2,943	26.7
Bingo-E (Ours)	80.6	779	76.7	86.7	345	84.9	36.7	1,584	33.0	33.3	2,943	26.7
<i>DeepSeek-R1-Distill-Qwen-7B</i>												
Base	82.8	3,033	65.7	85.7	1,001	80.3	37.8	4,340	25.9	40.0	6,528	18.0
PPO	88.4	1,536	79.7	92.9	918	87.5	45.4	2,709	37.1	56.7	5,857	30.3
Bingo-A (Ours)	88.8	1,400	80.9	92.3	371	90.2	45.2	1,908	39.6	63.3	4,670	41.5
Bingo-E (Ours)	87.2	1,252	80.3	91.8	366	89.7	45.0	1,693	40.1	60.0	4,011	42.9
<i>Qwen2.5-Math-7B-Instruct</i>												
Base	80.8	727	70.3	95.8	331	90.3	36.8	919	30.7	16.7	1,310	12.5
PPO	82.0	670	72.3	96.6	305	91.5	37.6	759	32.5	20.0	1,260	15.2
Bingo-A (Ours)	82.6	656	73.0	96.1	283	91.5	37.9	598	34.0	20.0	892	16.8
Bingo-E (Ours)	81.6	559	73.6	96.0	241	92.0	37.1	552	33.5	16.7	811	14.2

to Vanilla PPO (e.g., 894 vs. 2,771 tokens on MATH500), demonstrating the model’s ability to generate concise and correct reasoning steps.

Existing baselines struggle with the trade-off between accuracy and brevity: Approaches like Efficient Reasoning produce verbose outputs, while methods such as O1-Pruner overly shorten responses, compromising accuracy.

4.3 Performance Evaluation across Varying Model Scales

BINGO achieves the best trade-off between accuracy and response length across different model sizes: As shown in Table 2, both BINGO-A and BINGO-E consistently outperform all other methods across various model sizes (1.5B and 7B parameters) and benchmarks, achieving the highest L-Acc while maintaining competitive or superior accuracy.

BINGO-E offers a substantial reduction in

response length without sacrificing accuracy: BINGO-E reduces response length by up to 63% (e.g., 366 vs. 1,001 tokens on GSM8K) compared to the Base model, while also improving accuracy by 6.1 percentage points, demonstrating the model’s ability to generate concise and accurate reasoning steps.

4.4 Analysis of Significant versus Insignificant Token Ratio

Our significance-aware reward consistently increases the proportion of significant tokens across all datasets (e.g., from 0.71 to 0.75 on MATH500), leading to shorter and more focused reasoning chains, as illustrated in Figure 3. By reducing redundant tokens, this results in lower response length and improved Acc and L-Acc, demonstrating enhanced reasoning efficiency without sacrificing essential content.




Figure 3: **Significant Length Ratio dynamics during training.** The x-axis indicates training steps, and the y-axis denotes the proportion of significant tokens in the generated responses. Each subplot corresponds to one benchmark evaluated using DeepSeek-R1-Distill-Qwen-1.5B as the base model. The blue curve represents the baseline method (Vanilla PPO), and the red curve represents our approach (Ours).

Table 3: **Ablation study on reward components.** Each method is evaluated using DeepSeek-R1-Distill-Qwen-1.5B as the base model by answer accuracy (Acc, %), response length (Len), and length-normalized accuracy (L-Acc, %). Values in parentheses indicate the relative drop in L-Acc compared to BINGO-A. The best performance is highlighted in **dark blue**, and the second-best in **light blue**.

Method	MATH500			GSM8K			TheoremQA			AIME2024		
	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow
Bingo-A (Ours)	82.2	894	77.6	87.0	563	83.9	36.8	1,648	32.9	33.3	2,943	26.7
Vanilla PPO	81.4	2,771	66.2 (-14.7)	85.4	1,310	78.2 (-6.8)	32.3	4,146	22.7 (-31.0)	26.7	6,961	10.3 (-16.4)
Significance-Aware Length Reward	81.4	1,734	72.3 (-5.3)	86.7	742	82.6 (-1.3)	36.0	2,841	29.1 (-3.8)	40.0	5,138	24.4 (-2.3)
w/o Cosine	78.6	1,750	69.7 (-7.9)	85.7	509	83.0 (-0.9)	35.3	2,414	29.7 (-3.2)	33.3	6,454	15.4 (-11.3)
w/o Significance Separation	79.8	1,666	71.2 (-6.4)	86.6	604	83.3 (-0.6)	36.9	2,328	31.3 (-1.6)	26.7	5,702	14.7 (-12.0)
w/o Length Incentive	77.8	1,400	70.8 (-6.8)	82.6	425	80.5 (-3.4)	35.7	1,636	32.0 (-0.9)	30.0	4,157	21.1 (-5.6)
Dynamic Length Reward	79.0	2,204	67.5 (-10.1)	84.3	955	79.2 (-4.7)	33.9	2,632	27.9 (-5.0)	30.0	5,047	18.6 (-8.1)

4.5 Ablation Study

Combining Significance-Aware and Dynamic Length rewards yields the best trade-off: Table 3 shows that the joint use of both the Significance-Aware and Dynamic Length rewards (BINGO-A) provides the best performance, achieving the highest accuracy and L-Acc across all four benchmarks, while maintaining competitive or superior raw accuracy compared to other methods.

Removing key reward components degrades performance significantly: Ablations show that removing any of the reward components leads to noticeable performance drops, particularly in terms of L-Acc.

4.6 Additional Experiments and Analysis

BINGO improves across multiple RL algorithms: We evaluate the generalizability of our reward design by integrating it into other RL algorithms, including RLOO, GRPO, and Reinforce++. As shown in Appendix J, BINGO variants consistently outperform vanilla ones, achieving superior performance in both accuracy and L-Acc.

BINGO effectively reduces response length, especially for incorrect samples: The distribution of response lengths for correct vs. incorrect samples in Appendix N shows that BINGO significantly

shortens incorrect sample lengths compared to PPO. Furthermore, Figure 9 in Appendix N illustrates that incorrect samples show a more significant reduction in response length during training, confirming the dynamic reward’s effectiveness. Figure 8 in Appendix M further shows that our method consistently reduces response length more than PPO.

Analysis of significant token ratio and token-level significance visualization: Appendix K shows that BINGO increases the proportion of significant tokens compared to baselines, while Appendix L provides a token-level significance visualization, demonstrating how our approach retains essential reasoning steps and eliminates redundancy. **Analysis confirms the effectiveness of dynamic and significance rewards:** The analysis in Appendix O validates that our dynamic and significance rewards balance exploration and efficiency. A case study in Appendix P further demonstrates the practical impact of BINGO on reasoning efficiency.

5 Conclusion

In this work, we revisit reasoning efficiency in LLMs from a token-level perspective and propose BINGO, an RL framework that jointly optimizes accuracy and efficiency. By incorporating token

significance into reward design, BINGO selectively penalizes insignificant tokens while preserving essential reasoning, and uses a dynamic length reward to balance exploration and conciseness during training, achieving a favorable trade-off between correctness and efficiency.

Limitations

Despite the promising results achieved by BINGO, several aspects warrant further exploration. First, our current study instantiates BINGO within a PPO-based training framework, and we leave a more comprehensive investigation of its compatibility with other reinforcement learning algorithms to future work. While the underlying design is general, different optimization schemes may require tailored adaptations to fully realize its potential. Second, since BINGO follows standard PPO-style optimization, its training cost is similar to that of PPO. Future work may explore more efficient training strategies, such as curriculum learning or partial parameter updates, to further improve scalability.

Acknowledgments

Life, much like the realm of reasoning where complexity often clouds clarity, is filled with choices. Each decision, each token, carries with it uncertainty, hesitation, and fear. Yet amidst this ambiguity, we seek something deeper: to make each step joyful and meaningful, much like selecting only the most significant tokens. BINGO is our call to action, a declaration that precision and progress emerge not from waiting but from doing. We must summon the courage to take that first bold step, embracing uncertainty without fear and moving forward without hesitation. In doing so, we choose joy, and we choose meaning. As we journey onward, token by token and challenge by challenge, we find that with each thoughtful step, the path unfolds. And in the end, the answer becomes clear, resounding unmistakably: **Bingo**. Thanks for everything throughout this journey.

References

Pranjal Aggarwal and Sean Welleck. 2025. **L1: Controlling how long a reasoning model thinks with reinforcement learning**. *Preprint*, arXiv:2503.04697.

Daman Arora and Andrea Zanette. 2025. **Training language models to reason efficiently**. *Preprint*, arXiv:2502.04463.

Lang Cao. 2024. **GraphReason: Enhancing reasoning capabilities of large language models through a graph-based verification approach**. In *Proceedings of the 2nd Workshop on Natural Language Reasoning and Structured Explanations (@ACL 2024)*, pages 1–12, Bangkok, Thailand. Association for Computational Linguistics.

Lang Cao, Jingxian Xu, Hanbing Liu, Jinyu Wang, Mengyu Zhou, Haoyu Dong, Shi Han, and Dongmei Zhang. 2025. **Fortune: Formula-driven reinforcement learning for symbolic table reasoning in language models**. *arXiv preprint arXiv:2505.23667*.

Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. 2023. **Theoremqa: A theorem-driven question answering dataset**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7889–7901, Singapore. Association for Computational Linguistics.

Paul Francis Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. **Deep reinforcement learning from human preferences**. *ArXiv*, abs/1706.03741.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. **Training verifiers to solve math word problems**. *arXiv preprint arXiv:2110.14168*.

Ganqu Cui, Yuchen Zhang, Jiacheng Chen, Lifan Yuan, Zhi Wang, Yuxin Zuo, Haozhan Li, Yuchen Fan, Huayu Chen, Weize Chen, and 1 others. 2025. **The entropy mechanism of reinforcement learning for reasoning language models**. *arXiv preprint arXiv:2505.22617*.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. **Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning**. *Preprint*, arXiv:2501.12948.

Suriya Gunasekar, Yi Zhang, Jyoti Aneja, Caio César Teodoro Mendes, Allie Del Giorno, Sivakanth Gopi, Mojan Javaheripi, Piero Kauffmann, Gustavo de Rosa, Olli Saarikivi, Adil Salim, Shital Shah, Harkirat Singh Behl, Xin Wang, Sébastien Bubeck, Ronen Eldan, Adam Tauman Kalai, Yin Tat Lee, and Yuanzhi Li. 2023. **Textbooks are all you need**. *Preprint*, arXiv:2306.11644.

Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. 2025. **Token-budget-aware llm reasoning**. *Preprint*, arXiv:2412.18547.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. **Measuring mathematical prob-**

- lem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Le Hou, Richard Yuanzhe Pang, Tianyi Zhou, Yuexin Wu, Xinying Song, Xiaodan Song, and Denny Zhou. 2022. **Token dropping for efficient BERT pretraining**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3774–3784, Dublin, Ireland. Association for Computational Linguistics.
- Jian Hu. 2025. **Reinforce++: A simple and efficient approach for aligning large language models**. *ArXiv*, abs/2501.03262.
- Pengcheng Jiang, Jiacheng Lin, Lang Cao, R. Tian, S. Kang, Z. Wang, Jimeng Sun, and Jiawei Han. 2025. Deepretrieval: Hacking real search engines and retrievers with large language models via reinforcement learning. *arXiv preprint arXiv: 2503.00223*.
- Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. 1996. **Reinforcement learning: A survey**. *J. Artif. Intell. Res.*, 4:237–285.
- Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. 2024. **C3ot: Generating shorter chain-of-thought without compromising effectiveness**. *Preprint*, arXiv:2412.11664.
- Yucheng Li, Bo Dong, Chenghua Lin, and Frank Guerin. 2023. **Compressing context to enhance inference efficiency of large language models**. *Preprint*, arXiv:2310.06201.
- Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhijiang Guo, Le Song, and Cheng-Lin Liu. 2025. **From system 1 to system 2: A survey of reasoning large language models**. *Preprint*, arXiv:2502.17419.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. **Let’s verify step by step**. *Preprint*, arXiv:2305.20050.
- Zhenghao Lin, Zhibin Gou, Yeyun Gong, Xiao Liu, Yelong Shen, Ruochen Xu, Chen Lin, Yujiu Yang, Jian Jiao, Nan Duan, and Weizhu Chen. 2025. **Rho-1: Not all tokens are what you need**. *Preprint*, arXiv:2404.07965.
- Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. 2025. **O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning**. *Preprint*, arXiv:2501.12570.
- Wenjia Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. 2025. **Reasoning models can be effective without thinking**. *Preprint*, arXiv:2504.09858.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025. **s1: Simple test-time scaling**. *Preprint*, arXiv:2501.19393.
- OpenAI. 2024. **Gpt-4 technical report**. *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis Christiano, Jan Leike, and Ryan J. Lowe. 2022. **Training language models to follow instructions with human feedback**. *ArXiv*, abs/2203.02155.
- Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin, Victor Rühle, Yuqing Yang, Chin-Yew Lin, H. Vicky Zhao, Lili Qiu, and Dongmei Zhang. 2024. **Lmlingua-2: Data distillation for efficient and faithful task-agnostic prompt compression**. *Preprint*, arXiv:2403.12968.
- Xiaoye Qu, Yafu Li, Zhaochen Su, Weigao Sun, Jianhao Yan, Dongrui Liu, Ganqu Cui, Daizong Liu, Shuxian Liang, Junxian He, Peng Li, Wei Wei, Jing Shao, Chaochao Lu, Yue Zhang, Xian-Sheng Hua, Bowen Zhou, and Yu Cheng. 2025. **A survey of efficient reasoning for large reasoning models: Language, modality, and beyond**. *Preprint*, arXiv:2503.21614.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. **Proximal policy optimization algorithms**. *ArXiv*, abs/1707.06347.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024. **Deepseekmath: Pushing the limits of mathematical reasoning in open language models**. *Preprint*, arXiv:2402.03300.
- Shuming Shi, Jian Zhang, Yi Shen, Kai Wang, Shiguo Lian, Ning Wang, Wenjing Zhang, Jieyun Huang, and Jiangze Yan. 2025. **Dast: Difficulty-adaptive slow-thinking for large reasoning models**. *Preprint*, arXiv:2503.04472.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan J. Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. **Learning to summarize from human feedback**. *ArXiv*, abs/2009.01325.
- Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Hanjie Chen, and Xia Hu. 2025. **Stop overthinking: A survey on efficient reasoning for large language models**. *Preprint*, arXiv:2503.16419.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung,

- Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2022. **Challenging big-bench tasks and whether chain-of-thought can solve them**. *Preprint*, arXiv:2210.09261.
- Kimi Team, Angang Du, Bofei Gao, Bawei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, Chuning Tang, Congcong Wang, Dehao Zhang, Enming Yuan, Enzhe Lu, Fengxiang Tang, Flood Sung, Guangda Wei, Guokun Lai, and 75 others. 2025. **Kimi k1.5: Scaling reinforcement learning with llms**. *Preprint*, arXiv:2501.12599.
- Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. **Solving math word problems with process- and outcome-based feedback**. *Preprint*, arXiv:2211.14275.
- Hemish Veeraboina. 2023. **Aime problem set 1983-2024**.
- Rui Wang, Hongru Wang, Boyang Xue, Jianhui Pang, Shudong Liu, Yi Chen, Jiahao Qiu, Derek Fai Wong, Heng Ji, and Kam-Fai Wong. 2025. **Harnessing the reasoning economy: A survey of efficient reasoning for large language models**. *Preprint*, arXiv:2503.24377.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. **Self-consistency improves chain of thought reasoning in language models**. *Preprint*, arXiv:2203.11171.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. **Emergent abilities of large language models**. *Preprint*, arXiv:2206.07682.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. **Chain-of-thought prompting elicits reasoning in large language models**. *Preprint*, arXiv:2201.11903.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022b. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Yuyang Wu, Yifei Wang, Tianqi Du, Stefanie Jegelka, and Yisen Wang. 2025. **When more is less: Understanding chain-of-thought length in llms**. *Preprint*, arXiv:2502.07266.
- Heming Xia, Yongqi Li, Chak Tou Leong, Wenjie Wang, and Wenjie Li. 2025. **Tokenskip: Controllable chain-of-thought compression in llms**. *Preprint*, arXiv:2502.12067.
- Jingxian Xu, Mengyu Zhou, Weichang Liu, Hanbing Liu, Shi Han, and Dongmei Zhang. 2025a. **Twt: Thinking without tokens by habitual reasoning distillation with multi-teachers’ guidance**. *Preprint*, arXiv:2503.24198.
- Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. 2025b. **Chain of draft: Thinking faster by writing less**. *Preprint*, arXiv:2502.18600.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. **Tree of thoughts: Deliberate problem solving with large language models**. *Preprint*, arXiv:2305.10601.
- Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neubig, and Xiang Yue. 2025. **Demystifying long chain-of-thought reasoning in llms**. *Preprint*, arXiv:2502.03373.
- Jintian Zhang, Yuqi Zhu, Mengshu Sun, Yujie Luo, Shuofei Qiao, Lun Du, Da Zheng, Huajun Chen, and Ningyu Zhang. 2025. **Lightthinker: Thinking step-by-step compression**. *Preprint*, arXiv:2502.15589.

Contents of Appendix

A Ethical Considerations	13
B Reproducibility Statement	13
C Discussion of Token Significance Measurement	13
D Theoretical Analysis of Significance-Aware Length Reward	14
E Theoretical Discussion of Dynamic Length Reward	15
F Definition and Theoretical Analysis of Length-normalized Accuracy	16
G Detailed Settings of Experiments	17
H Performance under Extended Sampling Settings	19
I Performance overview of BINGO and other baselines	19
J Performance across Different Reinforcement Learning Algorithms	20
K Analysis of Significant Token Ratio	20
L Token-level Significance Visualization	20
M Analysis of Response Lengths Trends during Training	21
N Analysis of Response Lengths Dynamics for Correct vs. Wrong Samples	21
O Analysis of Incorrect Response Length under Different Reward Designs	23
P Case Study	23
Q Hyperparameter Study	23
R Notation Table	24

A Ethical Considerations

Our research focuses on improving the efficiency of large language model reasoning through reinforcement learning techniques, which poses no direct ethical concerns regarding human subjects, as no human data collection or experimentation was conducted. All datasets used (MATH, GSM8K, TheoremQA, AIME2024) are publicly available benchmarks with proper citations. We acknowledge the broader implications of more efficient LLM reasoning, including potential dual-use concerns, but emphasize that our contributions aim to reduce computational costs and environmental impact of AI systems. The research was conducted with academic integrity, and all authors have reviewed and agree with the content presented. There are no conflicts of interest to declare.

B Reproducibility Statement

To ensure reproducibility, we provide comprehensive implementation details throughout the paper. Section 3 describes our complete algorithmic framework, including all hyperparameters and reward formulations. Section 4.1 details our experimental setup, while Appendix G provides comprehensive settings including training and evaluation configurations, dataset settings, optimization parameters, data splits, evaluation metrics, and computational requirements. All experiments use publicly accessible pre-trained models (DeepSeek-R1-Distill-Qwen-1.5B, DeepSeek-R1-Distill-Qwen-7B, and Qwen2.5-Math-7B-Instruct) and datasets available on HuggingFace. Our code is available through an anonymous link in the abstract and as a zip file in the supplemental materials, and will be made publicly available upon acceptance.

C Discussion of Token Significance Measurement

As introduced in Section 3.1, token significance is computed using a contextual scoring function. In our work, this function is instantiated with *LLMLingua-2* (Pan et al., 2024). Given a generated sequence $y = (y_1, y_2, \dots, y_n)$, each token y_i is assigned a significance score

$$S(y_i) = P(y_i | \mathbf{y}_{\leq n}; \theta_{\mathcal{M}_e}), \quad (12)$$

where $\theta_{\mathcal{M}_e}$ denotes the parameters of a bidirectional encoder model. By conditioning on both preceding and succeeding context, this formulation

Table 4: Comparison between manual token significance annotations and *LLMLingua-2*-generated labels on 100 representative samples.

Metric	Score
Average Precision	0.8026
Average Recall	0.9222
Average F1-score	0.8518

enables a balanced assessment of token-level contribution within the full sequence and mitigates positional effects inherent to left-to-right estimation. In practice, *LLMLingua-2* provides a lightweight and computationally efficient instantiation of this contextual scoring function.

To further assess whether *LLMLingua-2* provides reliable token significance estimates for our setting, we manually annotated 100 representative samples and compared human labels against *LLMLingua-2*-generated labels. In the manual annotation, tokens labeled as *significant* mainly included key derivation steps, numerical values, formulas, and essential action verbs, while tokens labeled as *insignificant* mainly included common linking words, fillers, and conjunctions such as “and” and “or.” The comparison results are reported in Table 4. *LLMLingua-2* achieves an average precision of 0.8026, an average recall of 0.9222, and an average F1-score of 0.8518. In particular, the high recall indicates that *LLMLingua-2* successfully captures most tokens that humans identify as significant. Although minor discrepancies remain, which is expected given linguistic ambiguity and annotation subjectivity, the overall results suggest that *LLMLingua-2* provides sufficiently accurate and robust token significance estimates for our purpose.

An alternative approach is *Selective Context* (Li et al., 2023), which defines token importance using a unidirectional language modeling objective:

$$S(y_i) = -\log P(y_i | y_{<i}; \theta_{M_L}), \quad (13)$$

where θ_{M_L} denotes the parameters of a unidirectional language model. Under this formulation, token importance is interpreted through model uncertainty, with less predictable tokens receiving higher scores. However, because it relies exclusively on past context, this approach is prone to position bias and cannot capture interactions between a token and subsequent reasoning steps, thereby limiting its ability to assess token-level contribution within the complete reasoning trace.

In contrast to such unidirectional uncertainty-based measures, contextual scoring based on bidirectional encoders offers a more holistic view of token significance. By evaluating tokens with respect to the entire sequence, it better aligns with the objective of discouraging redundant or low-contribution tokens while preserving essential reasoning content. The manual evaluation above further supports that *LLMLingua-2* can identify human-recognized significant tokens with high fidelity, making it a suitable choice for our token significance measurement. The effectiveness of this type of contextual, bidirectional scoring has also been supported by recent work such as TokenSkip (Xia et al., 2025), which demonstrates its ability to identify redundant components in chain-of-thought reasoning.

D Theoretical Analysis of Significance-Aware Length Reward

Preliminaries and Notation. Given a prompt x , the policy π_θ generates a chain-of-thought (CoT) sequence $Y = (y_1, \dots, y_T)$, from which a deterministic decoder produces a final answer $\hat{Z} \in \mathcal{Z}$. To assess the relative informativeness of each token, we compute a *significance score* using *LLMLingua-2* (Pan et al., 2024):

$$S(y_i) = P(y_i | Y; \theta_{\mathcal{M}_e}), \quad (14)$$

where $\theta_{\mathcal{M}_e}$ denotes the parameters of a bidirectional encoder model. Unlike unidirectional predictors, *LLMLingua-2* uses both left and right context to provide a holistic estimate of token informativeness. Based on a fixed threshold τ , we partition the sequence as:

$$\mathcal{Y}_{\text{sig}} = \{y_i : S(y_i) \geq \tau\}, \quad (15)$$

$$\mathcal{Y}_{\text{insig}} = \{y_i : S(y_i) < \tau\}, \quad (16)$$

where \mathcal{Y}_{sig} and $\mathcal{Y}_{\text{insig}}$ denote the sets of *significant* and *insignificant* tokens, respectively.

Motivation for a Mutual Information Proxy. In principle, each token’s importance could be measured by its mutual information with the final answer, $I(y_i; Z^*)$. However, computing the exact joint distribution $p(y_i, Z^*)$ is intractable due to the vast generation space and limited supervision. Instead, we employ a proxy that is (i) efficient to compute for each token and (ii) monotonically correlated with $I(y_i; Z^*)$.

LLMLingua-2 (Pan et al., 2024) satisfies these requirements by training under an information bottleneck objective:

$$I(T; Y) - \beta I(T; Z^*), \quad (17)$$

where T is the retained subsequence. Tokens with low conditional probability typically carry little additional information about Z^* , while high-probability tokens preserve essential semantics.

Assumption 1 (Fidelity of the Mutual Information Proxy). There exist constants $c > \varepsilon > 0$ such that

$$I(y_i; Z^*) \leq \varepsilon \quad (\forall y_i \in \mathcal{Y}_{\text{insig}}), \quad (18)$$

$$I(y_j; Z^*) \geq c - \varepsilon \quad (\forall y_j \in \mathcal{Y}_{\text{sig}}). \quad (19)$$

Lemma 1 (Bounded Accuracy Loss). Let \hat{Z}_{full} denote the answer decoded from the full CoT, and \hat{Z}_{sig} the answer decoded after removing $\mathcal{Y}_{\text{insig}}$. Under Assumption 1, the increase in error probability is bounded:

$$\left| \Pr[\hat{Z}_{\text{sig}} \neq Z^*] - \Pr[\hat{Z}_{\text{full}} \neq Z^*] \right| \leq \varepsilon. \quad (20)$$

Proof. Let $Y = (y_1, \dots, y_T)$ and $Y_{\text{sig}} = Y \setminus \mathcal{Y}_{\text{insig}}$. By the chain rule:

$$I(Y; Z^*) = I(Y_{\text{sig}}; Z^*) + \sum_{y_i \in \mathcal{Y}_{\text{insig}}} I(y_i; Z^* | Y_{<i}), \quad (21)$$

and each term in the sum is at most ε . Therefore,

$$I(Y; Z^*) - I(Y_{\text{sig}}; Z^*) \leq T\varepsilon, \quad (22)$$

and by Fano’s inequality, this gap translates into an error increase of at most ε . \square

Definition 1 (General vs. Significance-Aware Length Reward). For a generated trace Y , define two reward functions:

$$R_{\text{len}}(Y) = \mathbb{1}[\hat{Z}(Y) = Z^*] - \lambda |Y|, \quad (23)$$

$$R_{\text{sig}}(Y) = \mathbb{1}[\hat{Z}(Y) = Z^*] - \lambda |\mathcal{Y}_{\text{insig}}|. \quad (24)$$

Here, R_{len} penalizes total length, while R_{sig} penalizes only insignificant tokens.

Theorem 1 (Benefit of the Significance-Aware Reward). Let π_θ be updated by a single PPO step using either reward, with the same coefficient $\lambda > 0$. If

$$\lambda > \frac{\varepsilon}{\mathbb{E}_{\pi_\theta}[|\mathcal{Y}_{\text{sig}}|]}, \quad (25)$$

then

$$\mathbb{E}_{\pi_\theta}[R_{\text{sig}}(Y)] > \mathbb{E}_{\pi_\theta}[R_{\text{len}}(Y)]. \quad (26)$$

Proof. Lemma 1 implies

$$\mathbb{E}[R_{\text{sig}}] - \mathbb{E}[R_{\text{len}}] = \lambda \mathbb{E}[|\mathcal{Y}_{\text{sig}}|] - \Delta_{\text{acc}}, \quad (27)$$

where

$$0 \leq \Delta_{\text{acc}} \leq \varepsilon. \quad (28)$$

Under the stated bound on λ , the difference is strictly positive. \square

Practical Implication. The significance-aware reward achieves the same or greater length reduction with provably smaller accuracy degradation than a general length reward. By selectively penalizing insignificant tokens, it still encourages conciseness while maintaining fidelity. With *LLMLingua-2* providing a fast proxy for token–answer informativeness, this reward design supports both principled and practical optimization for efficient reasoning.

E Theoretical Discussion of Dynamic Length Reward

We provide a theoretical discussion of the motivation for our dynamic length reward schedule by addressing three key questions:

1. Why does encouraging longer chains of thought (CoT) during early training help exploration?
2. Why does applying a fixed length penalty throughout training limit performance?
3. Why does dynamically flipping the reward from positive to negative upon convergence yield better accuracy–efficiency trade-offs?

1. Longer CoT Enables Richer Exploration.

Let $P_t(L)$ denote the model’s distribution over output length L at training step t . Define the expected accuracy given length L as

$$\text{Acc}(L) = \Pr[\hat{Z} = Z^* \mid L(Y) = L], \quad (29)$$

where \hat{Z} is the predicted answer and Z^* the ground-truth. Empirically, $\text{Acc}(L)$ follows a saturating “S-curve”:

$$\frac{d}{dL} \text{Acc}(L) > 0 \quad \text{for } L < L^*, \quad (30)$$

$$\frac{d}{dL} \text{Acc}(L) \approx 0 \quad \text{for } L \geq L^*, \quad (31)$$

where L^* is the length at which accuracy saturates. The expected accuracy at step t is

$$\text{Acc}_t = \sum_L P_t(L) \text{Acc}(L). \quad (32)$$

Shifting probability mass toward longer CoT (up to L^*) thus increases Acc_t , since longer reasoning expands exploration and raises the chance of discovering correct solution patterns.

Takeaway. Rewarding longer CoT early boosts exploration and accelerates convergence toward high accuracy.

2. Static Length Penalty Causes Premature Compression.

Consider a fixed penalty $\lambda > 0$, yielding reward

$$J_{\text{static}}(L) = \text{Acc}(L) - \lambda L. \quad (33)$$

The optimal length L_s under this objective satisfies

$$\left. \frac{d}{dL} \text{Acc}(L) \right|_{L=L_s} = \lambda. \quad (34)$$

Since $\frac{d}{dL} \text{Acc}(L)$ vanishes for $L \geq L^*$, any $\lambda > 0$ forces $L_s < L^*$, implying

$$\text{Acc}(L_s) < \text{Acc}(L^*). \quad (35)$$

Thus the model truncates its CoT before accuracy has fully converged.

Takeaway. Static penalties enforce efficiency too early, sacrificing potential accuracy gains.

3. Dynamic Penalty Supports a Two-Phase Curriculum.

We introduce a time-dependent penalty λ_t :

$$\lambda_t = \begin{cases} \gamma, & t < t_0 \quad (\text{exploration phase}), \\ \alpha(t - t_0), & t \geq t_0 \quad (\text{compression phase}), \end{cases} \quad (36)$$

where $\gamma < 0$, and t_0 is the step at which validation accuracy stabilizes, i.e.,

$$\Delta \text{Acc}_t = \frac{\text{Acc}_t - \text{Acc}_{t-\Delta}}{\Delta} < \beta. \quad (37)$$

Phase I (Exploration). During early training, we set $\lambda_t < 0$, effectively turning the penalty into a bonus:

$$J(L) = \text{Acc}(L) - \lambda_t L, \quad \text{with } -\lambda_t > 0,$$

which encourages longer outputs. Since $\text{Acc}(L)$ increases with L up to L^* , this promotes

$$L_t \rightarrow L^*, \quad \text{Acc}_t \rightarrow \text{Acc}(L^*). \quad (38)$$

Phase II (Compression). As training progresses, λ_t transitions from negative to positive and increases gradually. When $\lambda_t > 0$, the derivative of the reward at L^* is

$$\begin{aligned} \frac{d}{dL} [\text{Acc}(L) - \lambda_t L] \Big|_{L=L^*} &= \\ \frac{d}{dL} \text{Acc}(L^*) - \lambda_t &< 0, \end{aligned} \quad (39)$$

so extending beyond L^* reduces reward. The policy thus shortens to a new equilibrium L_d :

$$\frac{d}{dL} \text{Acc}(L) \Big|_{L=L_d} = \lambda_t, \quad L_d < L^*, \quad (40)$$

$$\text{Acc}(L_d) \approx \text{Acc}(L^*). \quad (41)$$

Comparison to Static Penalty. The final dynamic reward is

$$J_{\text{dyn}} = \text{Acc}(L_d) - \lambda_T L_d, \quad (42)$$

and under the concavity of $\text{Acc}(L)$ one can show

$$\begin{aligned} J_{\text{dyn}} - J_{\text{static}} &= [\text{Acc}(L_d) - \text{Acc}(L_s)] \\ &\quad - \lambda_T (L_d - L_s) \geq 0, \end{aligned} \quad (43)$$

i.e., dynamic scheduling yields no worse and often strictly better reward. This holds because for concave functions

$$\text{Acc}(L_d) - \text{Acc}(L_s) \geq \text{Acc}'(L_d)(L_d - L_s), \quad (44)$$

and with $\text{Acc}'(L_d) = \lambda_T$, the inequality follows.

Efficiency Metric. Define length-normalized accuracy

$$\text{L-Acc}(L) = \text{Acc}(L) \sqrt{1 - \frac{L}{L_{\max}}}. \quad (45)$$

In practice, dynamic scheduling often achieves similar or higher accuracy with shorter or comparable length, leading to

$$\text{L-Acc}(L_d) > \text{L-Acc}(L_s). \quad (46)$$

Conclusion. Our dynamic length reward realizes the curriculum

$$\begin{aligned} \text{explore freely } (\lambda \leq 0) &\longrightarrow \text{accuracy convergence} \\ &\longrightarrow \text{gradual compression } (\lambda \uparrow). \end{aligned} \quad (47)$$

This schedule lets the model reach its accuracy ceiling $\text{Acc}(L^*)$ before enforcing brevity, achieving better accuracy–efficiency trade-offs than static schemes.




Figure 4: **Penalty curve:** $\sqrt{1 - \frac{L}{L_{\max}}}$.

F Definition and Theoretical Analysis of Length-normalized Accuracy

Length-Normalized Accuracy.

To evaluate reasoning efficiency, we adopt a length-normalized accuracy metric, denoted as L-ACC, which balances correctness with brevity. Formally, it is defined as:

$$\text{L-Acc} = \text{Acc} \times \sqrt{1 - \frac{L}{L_{\max}}}, \quad (48)$$

where $\text{Acc} \in [0, 1]$ denotes exact-match accuracy, L is the number of tokens in the model’s response, and L_{\max} is a dataset-specific upper bound on response length. The multiplicative factor penalizes longer outputs in a sub-linear manner, rewarding models that solve problems with fewer tokens.

Specifically, we set $L_{\max} = 8192$ for the two *DeepSeek*-based reasoning models, and $L_{\max} = 3000$ for the *Qwen2.5-Math-7B-Instruct* model, since reasoning-oriented models generally generate longer outputs than instruction-tuned models. The multiplicative factor $\sqrt{1 - \frac{L}{L_{\max}}}$ weights accuracy by a sub-linear penalty on sequence length, so the metric rewards correct solutions that are delivered with fewer tokens. Normalizing by L_{\max} makes the score comparable across datasets of very different scale, while the square-root ensures a smooth, continuous trade-off: the first tokens cut away improve the score more than later ones, mirroring human tolerance for moderate verbosity but aversion to extreme length. When $L = L_{\max}$ the metric collapses to zero, preventing models from exchanging unbounded length for marginal accuracy gains; when $L = 0$ it reduces to the raw accuracy, preserving credit for perfectly concise answers.

Penalty Behavior and Physical Intuition. The penalty term $\sqrt{1 - \frac{L}{L_{\max}}}$ is continuous, monotonically decreasing, and bounded between 0 and 1. It applies no penalty when $L = 0$, and reduces the reward to zero when $L = L_{\max}$, even if the answer is correct. Crucially, the square-root form introduces diminishing returns: trimming early redundant tokens provides larger gains in L-ACC than removing tokens later in the sequence. This design mirrors human preferences—we tolerate moderate verbosity, but disfavor excessive detail. It also echoes the behavior of L2 regularization, where larger values are penalized more aggressively, while smaller deviations are softly constrained.

Gradient Analysis. To understand its optimization implications, we analyze the gradient of the penalty term with respect to L :

$$\frac{d}{dL} \left(\sqrt{1 - \frac{L}{L_{\max}}} \right) = -\frac{1}{2L_{\max}} \cdot \left(1 - \frac{L}{L_{\max}} \right)^{-1/2} \quad (49)$$

This derivative diverges as $L \rightarrow L_{\max}$, indicating that long outputs are heavily penalized. In contrast, when L is small, the gradient approaches zero, and the penalty becomes negligible. This behavior encourages models to first eliminate highly redundant tokens, while maintaining stability for shorter outputs.

Optimization Benefits. Unlike hard constraints on length, this formulation yields a smooth and differentiable reward signal, making it well-suited for reinforcement learning algorithms such as PPO. It provides stable guidance throughout training and enables the model to trade off between accuracy and length in a controlled and interpretable manner. As shown in Figure 4, the penalty curve strongly discourages excessively long outputs while allowing flexibility in moderately verbose cases, contributing to more efficient and human-aligned reasoning. Notably, the curve becomes steep as the response length approaches L_{\max} , meaning that small increases in length lead to sharp drops in reward; conversely, it flattens near $L = 0$, where changes in length have only a minor effect on the reward. This property ensures that the model is heavily penalized for extreme verbosity while remaining tolerant of brief explanatory content.

G Detailed Settings of Experiments

Prompt. All experiments use the prompt: "Let’s think step by step and output the final

answer within `\boxed{\}`."

Models. Our experiments involve a mix of proprietary and open-source models. The models evaluated in this study include:

- **DeepSeek-R1-Distill-Qwen-1.5B (MIT License):** A fine-tuned model with 1.5 billion parameters, used to evaluate the proposed method.
- **DeepSeek-R1-Distill-Qwen-7B (MIT License):** A fine-tuned model with 7 billion parameters, also used to evaluate the proposed method.
- **Qwen2.5-Math-7B-Instruct (Apache-2.0 License):** An instruction-tuned model with 7 billion parameters, used to further assess the efficiency and accuracy in reasoning tasks.

Datasets. We evaluate our models on several datasets covering both in-distribution (ID) and out-of-distribution (OOD) tasks. The evaluation framework encompasses:

- **MATH (Hendrycks et al., 2021):** A comprehensive training dataset containing 7,500 mathematical problems across various difficulty levels and topics.
- **MATH500 (Hendrycks et al., 2021):** A carefully selected 500-problem subset from the MATH test set, serving as our primary in-distribution evaluation benchmark.
- **GSM8K (Cobbe et al., 2021):** Grade school math word problems requiring multi-step reasoning, used for out-of-distribution evaluation on elementary-level mathematics.
- **TheoremQA (Chen et al., 2023):** A challenging dataset requiring theorem application and symbolic reasoning across STEM domains, used for out-of-distribution evaluation.
- **AIME2024 (Veeraboina, 2023):** Problems from the prestigious American Invitational Mathematics Examination, representing the most challenging out-of-distribution evaluation.

These datasets are arranged in increasing order of difficulty: $\text{GSM8K} < \text{MATH500} < \text{THEOREMQA} < \text{AIME2024}$, offering a comprehensive evaluation of models’ reasoning capabilities across varying complexity levels, as summarized in Table 5.

Preprocessing and Tokenization. Each model uses its corresponding tokenizer to process the input sequences. Tokenization ensures compatibility

Table 5: Overview of datasets used for training and evaluation.

Type	Dataset	# Train	# Test	Domain	Task Type	Difficulty	Source
Training	MATH (Hendrycks et al., 2021)	7,500	–	Mathematics	Problem Solving	Mixed	Link
ID Test	MATH500 (Hendrycks et al., 2021)	–	500	Mathematics	Problem Solving	Medium-Hard	Link
OOD Test	GSM8K (Cobbe et al., 2021)	–	1,319	Elementary Math	Word Problems	Easy	Link
	TheoremQA (Chen et al., 2023)	–	800	STEM	Theorem Application	Hard	Link
	AIME2024 (Veeraboina, 2023)	–	30	Competition Math	Advanced Problem Solving	Very Hard	Link

with the model’s input structure, using special tokens to denote the start and end of sequences.

Training Procedure. All models are trained for a total of 50 epochs using the Proximal Policy Optimization (PPO) algorithm, optimizing for both accuracy and efficiency. The actor and critic models are initialized with the same parameters, and training is conducted with the following hyperparameters: actor learning rate = 5×10^{-5} , critic learning rate = 1×10^{-6} , mini-batch size = 512, and KL-divergence coefficient = 0.001. Evaluation is performed periodically at every training step to monitor progress, and the best model checkpoints are selected for final testing.

Decoding Configurations. We conduct both training and evaluation under carefully controlled decoding settings. During training, we adopt sampling generation with temperature = 1.0, top_k = –1, and top_p = 1.0 to encourage exploration, following the default configuration of the VERL framework for comparability with prior work. A single response ($n = 1$) is generated per prompt, with the maximum prompt length capped at 1,024 tokens for efficiency. The maximum response length is set to 8,192 tokens for DeepSeek models and 3,000 tokens for Qwen-Math models.

For evaluation, we emphasize efficiency and stability by adopting greedy decoding, consistent with VERL defaults and prior studies (Yeo et al., 2025; Cui et al., 2025). Specifically, evaluation uses greedy decoding with temperature = 0, and one response per input ($n = 1$). The same maximum response lengths as in training are applied (8,192 for DeepSeek, 3,000 for Qwen-Math). We also conducted experiments under extended sampling configurations, with comprehensive results presented in Appendix H.

Evaluation Metrics. We evaluate the models using the following metrics:

- **Exact Match (EM):** Measures the proportion of exact matches between the generated output and the ground-truth answer.
- **Response Length (Len):** Measures the num-

ber of tokens in the output sequence.

- **Length-Normalized Accuracy (L-Acc):** A metric that balances accuracy and efficiency by considering both correctness and response length.

Baselines. We compare the BINGO framework with the following baselines:

- **DAST (Shi et al., 2025):** Uses dynamic length penalties based on problem difficulty.
- **Efficient Reasoning (Arora and Zanette, 2025):** Scales down positive rewards to encourage brevity.
- **Kimi-k1.5 (Team et al., 2025):** Applies online length penalties.
- **O1-Pruner (Luo et al., 2025):** Applies offline length penalties based on length comparisons with reference sequences.
- **Demystifying (Yeo et al., 2025):** Applies a symmetric penalty strategy for response lengths, encouraging both shorter and more extensive reasoning depending on correctness.

Rather than directly comparing published baselines—which employ diverse frameworks (e.g., SimPO, GRPO) and differ in their on-policy versus off-policy implementations—we isolate and re-implement only the length-based reward components proposed in each work. All these reward designs are integrated into a unified PPO framework. This approach enables a fair comparison focused specifically on the effectiveness of different reward formulations for improving reasoning efficiency.

Software and Hardware. The experiments are conducted with Python 3.11, PyTorch v2.4.0, and CUDA 12.8 for model training and inference. We use 4 NVIDIA A100 80GB PCIe GPUs for training the 7B model and 2 NVIDIA H100 80GB PCIe GPUs for training the 1.5B model. For inference, 2 NVIDIA H100 80GB PCIe GPUs are used to accelerate processing.

Table 6: **Performance comparison under sampling decoding settings.** Each method is evaluated using DeepSeek-R1-Distill-Qwen-1.5B as the base model with sampling parameters (32,768 token limit, 3 samples, temperature = 0.6, top-p = 1.0). Metrics include answer accuracy (Acc, %), response length (Len), and length-normalized accuracy (L-Acc, %). The best performance is highlighted in **dark blue**, and the second-best in **light blue**.

Method	MATH500			GSM8K			TheoremQA			AIME2024		
	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow
Base	81.6	5,155	74.9	83.7	1,748	81.4	31.7	7,598	27.8	17.8	15,703	12.8
Vanilla PPO	82.3	2,694	78.8	86.5	1,050	85.1	33.4	3,616	31.5	28.9	7,389	25.4
O1-Pruner	80.1	1,283	78.5	85.2	352	84.7	34.7	1,095	34.1	28.9	4,636	26.8
Demystifying	81.3	1,945	78.9	86.8	483	86.2	35.2	1,863	34.2	30.0	5,891	27.2
DAST	83.5	2,053	80.8	84.1	375	83.6	35.4	2,954	33.8	36.7	5,072	33.7
<i>Bingo (Ours)</i>												
Bingo-A	85.1	1,114	83.6	88.4	483	87.7	37.9	1,592	37.0	38.9	3,110	37.0
Bingo-E	84.2	983	82.9	88.1	217	87.8	37.7	1,004	37.1	37.6	2,817	35.9

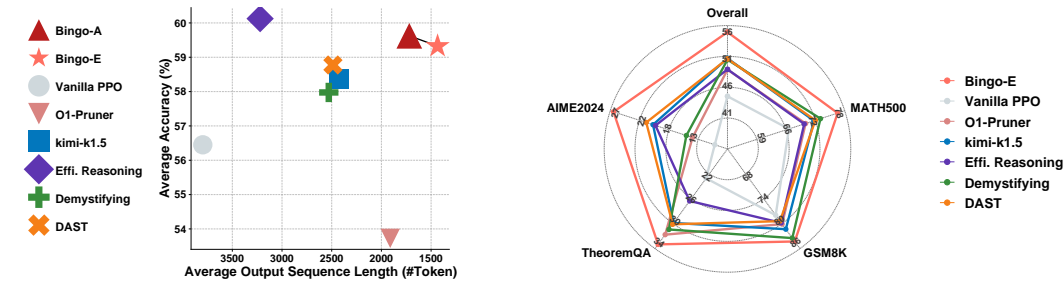


Figure 5: **Performance overview of BINGO and other baselines.** *Left:* Scatter plot of average accuracy versus average response length on four benchmarks (MATH500, GSM8K, TheoremQA, AIME2024) using DeepSeek-R1-Distill-Qwen-1.5B as the base model. Points nearer the top-right corner represent a better balance of accuracy and efficiency. *Right:* Radar chart of length-normalized accuracy for each method. Greater radial distances denote higher efficiency.

H Performance under Extended Sampling Settings

We conducted additional experiments using sampling decoding to assess the robustness of our approach under more exploratory conditions. These experiments employed an extended configuration with a 32,768-token output limit, three samples per prompt, temperature of 0.6, and top-p of 1.0. We evaluated the base DeepSeek-R1-Distill-Qwen-1.5B model, vanilla PPO, our proposed Bingo method, and selected competitive baselines to ensure comprehensive comparison.

Table 6 presents the accuracy and average response length across four benchmarks under these sampling conditions. The results demonstrate that Bingo maintains its efficiency advantage even with sampling decoding, achieving strong accuracy while generating substantially shorter outputs than all baseline methods. This finding confirms that our reward design effectively promotes concise reasoning regardless of the decoding strategy employed.

Although these extended settings yielded accu-

racy improvements, they required approximately five times the computational resources and training time compared to greedy decoding. Given this substantial computational overhead, we selected single-response greedy decoding as our primary evaluation protocol to maintain experimental feasibility while still providing meaningful performance assessments. The sampling results presented here validate that our approach remains effective under more computationally intensive conditions.

I Performance overview of BINGO and other baselines

Figure 5 provides an overview of the performance of BINGO compared with strong baselines in terms of both accuracy and efficiency. The scatter plot shows that BINGO variants consistently occupy the favorable region with higher accuracy and substantially shorter response lengths, indicating a superior accuracy–efficiency trade-off. This advantage holds across both relatively simple benchmarks such as GSM8K and more challenging settings like TheoremQA, where BINGO achieves noticeable

accuracy improvements while reducing response length by more than half. The radar chart further confirms these trends, with BINGO attaining the highest length-normalized accuracy across datasets, demonstrating its ability to generate concise yet effective reasoning across diverse difficulty levels.

J Performance across Different Reinforcement Learning Algorithms

To evaluate the generalizability of our reward design, we integrate it into multiple on-policy reinforcement learning (RL) algorithms, including PPO, RLOO, GRPO, and Reinforce++. As shown in Table 7, Bingo-enhanced variants consistently outperform their vanilla counterparts across all four benchmarks in both accuracy and length-normalized accuracy (L-Acc). Among them, BINGO-PPO delivers the strongest overall performance, achieving the highest or second-highest scores on all datasets while substantially reducing output length. Moreover, the benefits of Bingo extend beyond PPO: each Bingo variant improves L-Acc over its baseline by a clear margin, demonstrating that our reward formulation generalizes well across different policy optimization strategies. These gains are observed not only on in-distribution datasets like MATH500 and GSM8K, but also on more challenging out-of-distribution settings such as AIME2024, highlighting the robustness of our approach. Overall, the results confirm that a principled and learnable length-aware reward offers a consistent advantage across a variety of reasoning tasks and RL algorithms.

K Analysis of Significant Token Ratio

We trained DeepSeek-R1-Distill-Qwen-1.5B exclusively on the MATH corpus. To evaluate its generalization beyond the training distribution, we test the model on the in-distribution split MATH500 as well as three out-of-distribution (OOD) benchmarks: GSM8K, THEOREMQA, and AIME2024. Figure 6 shows that our approaches, **Bingo-A** and **Bingo-E**, achieve the most favorable length-accuracy trade-off across all four benchmarks.

- **Efficiency at peak accuracy.** At the checkpoints that obtain their highest accuracy, both Bingo variants require only about 20% of the tokens used by the *Base* model on MATH500, with similarly large reductions on GSM8K, THEOREMQA, and AIME2024.

- **Preservation of informative content.** Bingo increases the share of significant tokens to 75–81%, showing that the shortened rationales shed mainly redundant rather than essential reasoning steps.
- **Difficulty-dependent length trends.** Token counts grow with task difficulty: the two harder benchmarks, THEOREMQA and AIME2024, demand considerably longer rationales and yield lower absolute accuracy than MATH500 and GSM8K. Even under these tougher conditions, Bingo still delivers the highest accuracy while generating the fewest tokens.
- **Alleviating the length–accuracy trade-off.** Baselines that compress reasoning without accounting for token importance (e.g., *O1-Pruner*) exhibit marked accuracy declines, whereas Bingo maintains—and in some cases slightly improves—task performance.
- **Robustness across tasks.** The same advantage holds for algebraic, commonsense, formal-logic, and competition-style benchmarks, underscoring the generality of the significance-aware and dynamic length rewards.

These findings confirm that explicitly modeling token significance and adaptively scheduling length rewards enables language models to reason both *accurately* and *efficiently*.

L Token-level Significance Visualization

Figure 7 provides a token-level significance visualization for a sample reasoning task. The problem involves calculating the cost of yogurt based on a given sale, and both the BINGO and Vanilla PPO models generate responses to solve the problem. Each token in the generated response is color-coded based on its predicted significance, with red indicating significant tokens and blue representing insignificant ones. Darker shades of red and blue correspond to higher significance levels.

The response from BINGO (top) is notably shorter and more focused on the key reasoning steps, highlighting the model’s ability to concentrate on relevant tokens while avoiding unnecessary elaboration. In contrast, the Vanilla PPO response (bottom) is longer, with a higher proportion of redundant and less informative tokens, reflecting a less efficient reasoning process. This visualization clearly demonstrates how BINGO encourages more concise and targeted reasoning, optimizing for both

Table 7: **Comparison of reinforcement learning algorithms on four reasoning benchmarks.** Each method is evaluated using DeepSeek-R1-Distill-Qwen-1.5B as the base model by answer accuracy (Acc, %), response length (Len), and length-normalized accuracy (L-Acc, %). Bingo-based variants consistently outperform their vanilla counterparts across different RL optimizers (PPO, RLOO, GRPO, Reinforce++). Numbers in parentheses show the L-Acc gain over the corresponding vanilla baseline, with **green** indicating improvement.

Method	MATH500			GSM8K			TheoremQA			AIME2024		
	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow	Acc \uparrow	Len \downarrow	L-Acc \uparrow
Base	63.2	3,913	45.7	73.2	2,025	63.5	18.7	5,741	10.3	16.7	7,027	6.3
Vanilla PPO	81.4	2,771	66.2	85.4	1,310	78.2	32.3	4,146	22.7	26.7	6,961	10.3
Bingo-PPO	82.2	894	77.6 (+11.4)	87.0	563	83.9 (+5.7)	36.8	1,648	32.9 (+10.2)	33.3	2,943	26.7 (+16.4)
Vanilla RLOO	76.8	2,413	64.5	77.3	1,588	69.4	30.0	3,162	23.5	26.7	6,025	13.7
Bingo-RLOO	78.0	1,985	67.9 (+3.4)	80.7	450	78.5 (+9.1)	32.0	2,230	27.3 (+3.8)	33.3	5,583	18.8 (+5.1)
Vanilla GRPO	76.4	2,533	63.5	77.8	804	73.9	29.2	2,946	23.4	26.7	6,096	13.5
Bingo-GRPO	79.4	1,753	70.4 (+6.9)	80.0	449	77.8 (+3.9)	31.9	2,298	27.0 (+3.6)	30.0	5,886	15.9 (+2.4)
Vanilla Reinforce++	76.2	2,842	61.6	82.0	1,291	75.2	28.0	3,977	20.1	30.0	6,168	14.9
Bingo-Reinforce++	78.4	2,070	67.8 (+6.2)	81.0	640	77.8 (+2.6)	33.1	2,566	27.4 (+7.3)	30.0	5,885	15.9 (+1.0)




Figure 6: **Length-accuracy results for nine optimization algorithms on four datasets.** Bars show the number of tokens generated using DeepSeek-R1-Distill-Qwen-1.5B as the base model at the checkpoint that yields the reported accuracy (left axis). Each bar is partitioned into significant (dark) and insignificant (light) segments, and the percentage above the bar indicates the share of significant tokens. The solid line (right axis) gives the corresponding answer accuracy. Our methods, Bingo-A and Bingo-E, attain the highest accuracy while using the fewest tokens and achieving the greatest proportion of significant tokens, highlighting their superior reasoning efficiency.

accuracy and efficiency by emphasizing significant steps in the reasoning process.

M Analysis of Response Lengths Trends during Training

Figure 8 presents the evolution of response length over training steps for Vanilla PPO and our method on four benchmarks. Our approach consistently yields substantially shorter outputs than Vanilla PPO throughout training, demonstrating effective removal of redundant tokens, and converges more smoothly, reflecting robust length regularization. The reduction in response length is most pronounced on the more demanding tasks—MATH500 and AIME2024—where

Vanilla PPO produces very long sequences, yet our method maintains a compact reasoning footprint. Importantly, this improvement generalizes across diverse reasoning styles, from arithmetic problems in GSM8K and formal-logic questions in THEOREMQA to academic and competition-style challenges, confirming that our reward design enhances reasoning efficiency without compromising training stability.

N Analysis of Response Lengths Dynamics for Correct vs. Wrong Samples

Figure 9 illustrates how response length evolves for correct and incorrect samples under our approach.

Problem: Terry eats 2 yogurts a day. They are currently on sale at 4 yogurts for \$5.00. How much does he spend on yogurt over 30 days?

Response of BINGO:

Okay, so Terry eats 2 yogurt each day, and they're on sale where 4 yogurt cost \$5.00. I need to figure out how much he spend over 30 days. Hmm, let me break this down step by step. ... Since 4 yogurt cost \$5, and he needs 60 yogurt, which is 15 sets, so 15 times \$5 is \$75. Terry spends \$75.00 dollars on yogurt over 30 days.

(553 tokens with 57% Significant)

Response of Vanilla PPO:

Okay, so Terry eats 2 yogurt per day. He's planning to buy them on sale where they're 4 for \$5.00. I need to figure out how much he'll spend over 30 days. Wait, let me double-check that. ... Wait, another way to think about it: ... so $12 * \$5 = \60 . 52 yogurt is 13 sets, so $13 * \$5 = \65 . 56 yogurt is 14 sets, so $14 * \$5 = \70 . 60 yogurt is 15 sets, so $15 * \$5 = \75 . Yep, that's consistent. I think that's solid. So, Terry spends \$75.00 on yogurt over 30 days. Terry spends 75 dollars on yogurt over 30 days.

(916 tokens with 52% Significant)

Final Answer: 75.

Figure 7: **Token-level significance visualization for a sample reasoning task.** Each token is colored based on its predicted significance: red indicates significant tokens (darker = more significant), and blue indicates insignificant tokens (darker = less significant). The response from BINGO (top) is shorter and more concentrated around meaningful reasoning steps, while the Vanilla PPO response (bottom) is longer and contains more exploratory and redundant language. The visualization illustrates how Bingo encourages more efficient and focused reasoning.




Figure 8: **Response length trends during training across four datasets.** The y-axis shows the number of tokens generated per response using DeepSeek-R1-Distill-Qwen-1.5B as the base model; the x-axis denotes training steps. The red line represents our method, and the blue line corresponds to Vanilla PPO. Across all tasks, our method consistently produces shorter and more stable responses, demonstrating improved reasoning efficiency without compromising task performance.

In the early phase of training, incorrect samples produce markedly longer outputs than correct ones, demonstrating the impact of our length-incentive mechanism in promoting thorough exploration on challenging cases. As the dynamic length reward takes effect around mid-training, the length for wrong samples declines steeply—outpacing the reduction seen for correct samples—and the gap between the two curves narrows. By later stages, both curves converge toward similarly concise rationales, indicating that the model has learned to apply efficient reasoning uniformly. This behavior confirms that our combination of significance-aware and dynamic rewards not only drives exploration where needed but also enforces brevity once sufficient understanding is achieved, resulting in a balanced, adaptive pruning of redundant tokens.

swear correctness, we compare output length distributions of our method and the Vanilla PPO baseline across four benchmarks using *DeepSeek-R1-Distill-Qwen-1.5B*. As shown in Figure 10, correct responses consistently exhibit shorter lengths than incorrect ones across all tasks. Our method further produces sharply concentrated distributions for correct samples, suggesting more focused and efficient reasoning. In contrast, Vanilla PPO outputs are generally longer and more dispersed, with substantial overlap between correct and incorrect cases. Notably, the length of incorrect samples is substantially reduced compared to Vanilla PPO, suggesting that the dynamic reward mechanism—which gradually penalizes verbosity during training—plays a role in guiding more efficient responses.

To examine how response length relates to an-




Figure 9: **Response length dynamics for correct vs. wrong samples during training.** The x-axis indicates training steps, and the y-axis denotes response length in tokens for models trained on DeepSeek-R1-Distill-Qwen-1.5B as the base model. The blue line tracks correctly answered samples, while the yellow line tracks incorrectly answered samples. In the early stages, incorrect samples produce substantially longer responses, reflecting the effect of our length-incentive mechanism. As the dynamic length reward gradually diminishes, the response length for incorrect samples falls more sharply than that for correct samples, illustrating the model’s adaptive pruning of redundant reasoning steps.

O Analysis of Incorrect Response Length under Different Reward Designs

To gain deeper insights into the effect of our reward design during training, we analyze the significant response length of incorrect predictions—a proxy for the model’s exploratory behavior under uncertainty. As shown in Figure 11, our full method generates longer responses for incorrect samples compared to the version **without incentive** across all four benchmarks, indicating that the significance-aware reward successfully promotes deeper reasoning attempts when the model is uncertain. In contrast, removing the dynamic reward leads to consistently longer outputs throughout training, while our method gradually reduces response length over time. This trend confirms that dynamic reward scheduling effectively encourages concise reasoning as the model matures. Together, these findings highlight the complementary roles of the two reward components: one enhances exploration during failure, and the other improves efficiency during learning progression.

P Case Study

To further assess the reasoning quality and output efficiency of our method, we conduct two case studies comparing the Base model, PPO, and our BINGO framework across two model scales. As shown in Figures 12 and 13, the Base model often produces excessively long and repetitive outputs, sometimes exceeding 5,000 words, and either fails to converge or hides the correct answer deep within verbose reasoning or terminated generations without a final answer. PPO improves conciseness but occasionally generates incorrect outputs or re-

tains unnecessary repetition in intermediate steps. In contrast, BINGO consistently produces correct, well-structured solutions while significantly reducing output length—achieving up to 3x compression over PPO in the 7B setting. These examples demonstrate the effectiveness of our reward design in balancing reasoning depth and brevity, encouraging the model to generate focused and efficient reasoning even on complex tasks.

Q Hyperparameter Study

We evaluated several combinations of hyperparameters for BINGO on the GSM8K dataset using DeepSeek-R1-Distill-Qwen-1.5B as the base model. Table 8 reports the accuracy and output length across different settings.

Table 8: Performance of BINGO under different hyperparameter settings on GSM8K.

λ_c	λ_w^{is}	λ_w^s	S	β	α	τ	Acc.	Len.
2	2	5	5	2	0.5	0.5	86.6	570
2	2	5	10	5	0.2	0.8	86.7	585
5	5	5	10	2.5	0.4	0.6	86.9	578
5	5	5	10	2.5	0.5	0.5	87.0	563

Hyperparameter Definitions:

- λ_c : Insignificant Length Reward Weight for Correct Samples.
- λ_w^{is} : Insignificant Length Reward Weight for Incorrect Samples.
- λ_w^s : Significant Length Reward Weight for Incorrect Samples.
- S : Slope interval for the Dynamic Length Reward.

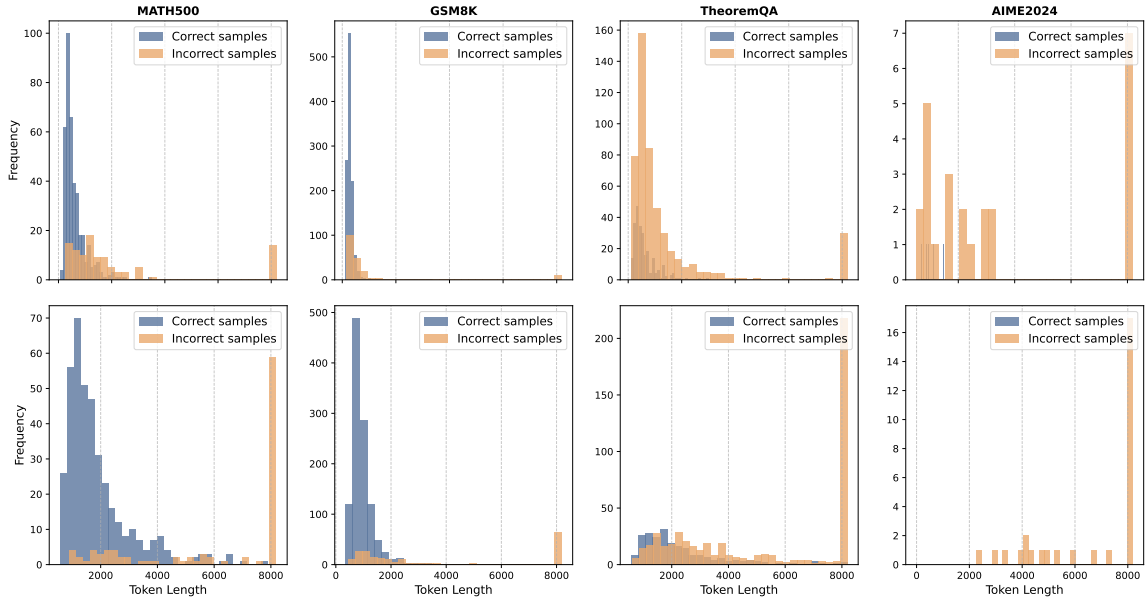


Figure 10: **Distribution of response lengths for correct vs. incorrect samples.** Histograms show the frequency of token lengths in model outputs across four benchmarks using DeepSeek-R1-Distill-Qwen-1.5B as the base model. Each plot compares correct responses (blue) and incorrect responses (orange). The top row corresponds to our method, while the bottom row shows results from Vanilla PPO. Across all datasets, incorrect samples are more likely to produce longer outputs, while correct samples tend to cluster in shorter length ranges. Compared to Vanilla PPO, our method produces a sharper, more compact distribution concentrated in shorter length regions.

- β : Threshold for Training Phase Transition.
- α : Decay Factor for Dynamic Length Reward.
- τ : Threshold for Significant Tokens.

As shown in Table 8, the performance of BINGO remains stable, with both accuracy and output length exhibiting only minor fluctuations across the tested hyperparameter ranges. This indicates that the method is robust to hyperparameter choices. Since the last configuration achieves the best overall performance, we fixed these hyperparameters for methods and datasets to ensure consistency and fairness in comparison.

R Notation Table

Table 9 offers a detailed overview of the notations utilized in this paper, along with their respective explanations. It serves as a handy reference to assist readers in grasping the concepts discussed in our work.

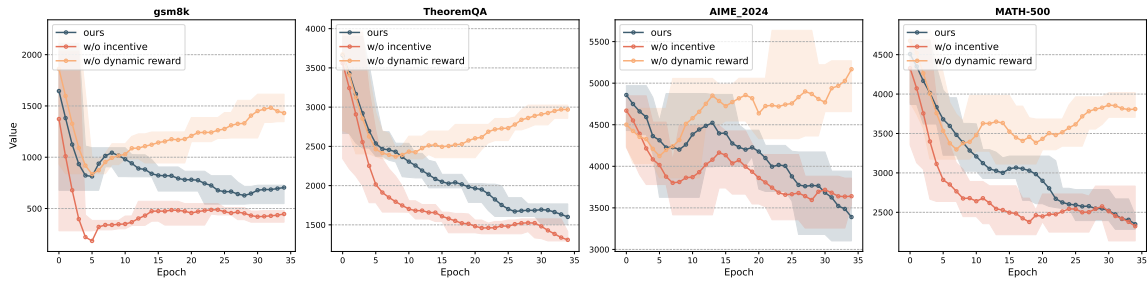


Figure 11: **Effect of Reward Design on Incorrect Response Length.** We visualize the average significant response length of incorrect predictions during training on four benchmarks using DeepSeek-R1-Distill-Qwen-1.5B as the base model. Compared to the variant without incentive, our full method produces longer responses for incorrect samples, suggesting that the significance-aware reward encourages more thorough exploration when the model is uncertain. In contrast, removing the dynamic reward leads to persistently longer outputs, whereas our full method shows a clear reduction in response length over time, confirming the effectiveness of dynamic reward scheduling in promoting concise reasoning. Together, these trends highlight the complementary roles of the two reward components in balancing exploration and efficiency.

Problem: Find the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$. (Do not use a calculator!)

✗ Response of Base Model with 2509 words:

Okay, so I need to find the greatest integer less than $((\sqrt{7} + \sqrt{5})^6/4)$. Hmm, that sounds a bit intimidating at first, but maybe I can break it down step by step. Let me think about how to approach this without a calculator, so I have to rely on algebraic manipulations and maybe some patterns or properties of exponents. First, I remember that expressions like $((\sqrt{a} + \sqrt{b})^n)$ can sometimes be simplified using binomial expansions, but that might get messy for the sixth power. Alternatively, I recall that sometimes pairing such terms can lead to simplifications. Maybe I can consider the conjugate of $(\sqrt{7} + \sqrt{5})$, which is $(\sqrt{7} - \sqrt{5})$, and see if their product is a nice number. Let me compute $(\sqrt{7} + \sqrt{5})(\sqrt{7} - \sqrt{5}) = 7 - 5 = 2$. Maybe I can find a recurrence relation or some pattern for $(x^n + y^n)$. Let me try computing some lower powers to see if I can spot a pattern. Let's compute $(x + y) : (x + y) = (\sqrt{7} + \sqrt{5}) + (\sqrt{7} - \sqrt{5}) = 2\sqrt{7}$. Hmm, that's interesting. Let's compute $(x^2 + y^2) : (x^2 + y^2) = 24$. Let's compute $(x^3 + y^3)$. Maybe I can use the identity $(x^3 + y^3) = (x + y)(x^2 - xy + y^2)$. But when I computed (x^6) directly, I got approximately 13535.988656. So, which is correct? Wait, perhaps I made a mistake in the direct expansion. Let me check the expansion again. Wait, in the direct expansion, I had: $(x^6 = 6768 + 1144\sqrt{35})$. So, $1144 \times \sqrt{35} \approx 1144 \times 5.916074 \approx 6767.988656$. So, $(x^6 = 6768 + 6767.988656)$. So, which is correct? Wait, perhaps I made a mistake in the direct expansion. Let me check the expansion step by step. So, that's correct. But in the direct expansion, I have $(x^6 \approx 13535.988674)$. So, which is correct? Wait, perhaps I made a mistake in the direct expansion. Let me check the expansion again. Wait, ...

✗ Response of PPO with 403 words:

To find the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$, we can use the property that $(\sqrt{7} + \sqrt{5})^n + (\sqrt{7} - \sqrt{5})^n$ is an integer for any positive integer n . Since $\sqrt{7} - \sqrt{5}$ is a small positive number (approximately 0.4142), $(\sqrt{7} - \sqrt{5})^6$ will be a small number close to 0. Let's denote $x = \sqrt{7} + \sqrt{5}$ and $y = \sqrt{7} - \sqrt{5}$. Then, $x + y = 2\sqrt{7}$ and $x - y = 2\sqrt{5}$. However, directly using this might not help. Instead, we know that $x^6 + y^6$ is an integer. Since $y < 1$, y^6 is very small, so x^6 is slightly less than an integer. Therefore, the greatest integer less than x^6 is $x^6 + y^6 - 1$. We need to compute $x^6 + y^6$. To find this, we can use recurrence relations or expand step by step. However, a more efficient way is to compute powers step by step: First, compute $x^2 = (\sqrt{7} + \sqrt{5})^2 = 7 + 2\sqrt{35} + 5 = 12 + 2\sqrt{35}$. Next, $y^2 = (\sqrt{7} - \sqrt{5})^2 = 7 - 2\sqrt{35} + 5 = 12 - 2\sqrt{35}$. Then, $x^3 + y^3 = (x + y)(x^2 + y^2 - xy) = 2\sqrt{7} \cdot (24 - 5) = 2\sqrt{7} \cdot 19 = 38\sqrt{7}$. Similarly, $x^3 - y^3 = (x - y)(x^2 + y^2 + xy) = 2\sqrt{5} \cdot (24 + 5) = 2\sqrt{5} \cdot 29 = 58\sqrt{5}$. However, this is getting complex. Instead, we can compute $x^6 + y^6$ using the recurrence relation. Let $S_n = x^n + y^n$. Then, $S_0 = 2$, $S_1 = 2\sqrt{7}$, $S_2 = 24$, $S_3 = 38\sqrt{7}$, $S_4 = 24 \cdot S_2 - S_1^2 = 24 \cdot 24 - (2\sqrt{7})^2 = 576 - 28 = 548$, $S_5 = S_4 \cdot S_1 - S_3 \cdot S_2 = 548 \cdot 2\sqrt{7} - 38\sqrt{7} \cdot 24 = (1096\sqrt{7} - 912\sqrt{7}) = 184\sqrt{7}$, $S_6 = S_5 \cdot S_1 - S_4 \cdot S_2 = 184\sqrt{7} \cdot 2\sqrt{7} - 548 \cdot 24 = 640 \cdot 2 - 548 \cdot 24 = 1280 - 13152 = -11972$. Thus, $x^6 + y^6 = -11972$. Since y^6 is positive, $x^6 = -11972 - y^6$. Therefore, x^6 is slightly less than -11972. But this contradicts since $x > 1$. Wait, actually, $y = \sqrt{7} - \sqrt{5} \approx 0.4142$, so $y^6 \approx (0.4142)^6 \approx 0.004$. Thus, $x^6 = -11972 - y^6 \approx -11972.004$, which is negative. This [Generation Termination without final answer]

✓ Response of Bingo with 470 words:

Okay, so I need to find the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$. Hmm, that sounds like I need to compute $(\sqrt{7} + \sqrt{5})^6$ and then take the floor of that value. But since I can't use a calculator, I have to figure out a way to compute this without approximating. I remember that expressions like $(\sqrt{a} + \sqrt{b})^n$ can sometimes be simplified using binomial expansion or by recognizing patterns. Maybe I can pair $(\sqrt{7} + \sqrt{5})$ with $(\sqrt{7} - \sqrt{5})$ because their product is $7 - 5 = 2$, which is nice. Let me denote $x = \sqrt{7} + \sqrt{5}$ and $y = \sqrt{7} - \sqrt{5}$. Then, $x \cdot y = 2$. I also know that $x + y = 2\sqrt{7}$ and $x - y = 2\sqrt{5}$. Since x and y are conjugates, maybe I can find a recurrence relation for $x^n + y^n$. Let me compute $x^1 + y^1 = 2\sqrt{7} \approx 5.2915$, $x^2 + y^2 = (x + y)^2 - 2xy = (2\sqrt{7})^2 - 2 \times 2 = 28 - 4 = 24$, $x^3 + y^3 = (x + y)(x^2 + y^2 - xy) = 2\sqrt{7} \times (24 - 2) = 2\sqrt{7} \times 22 = 44\sqrt{7} \approx 44 \times 2.6458 \approx 116.315$, $x^4 + y^4 = (x^2 + y^2)^2 - 2(xy)^2 = 24^2 - 2 \times (4) = 576 - 8 = 568$, $x^5 + y^5 = (x + y)(x^4 + y^4) - xy(x^3 + y^3) = 2\sqrt{7} \times 568 - 2 \times (44\sqrt{7}) = 1136\sqrt{7} - 88\sqrt{7} = 1048\sqrt{7} \approx 1048 \times 2.6458 \approx 2767.318$, $x^6 + y^6 = (x^3 + y^3)^2 - 2(xy)^3 = (44\sqrt{7})^2 - 2 \times (8) = 1936 \times 7 - 16 = 13552 - 16 = 13536$. So $x^6 + y^6 = 13536$. But since $y = \sqrt{7} - \sqrt{5}$, which is less than 1 (since $\sqrt{7} \approx 2.6458$, $\sqrt{5} \approx 2.236$, so $y \approx 0.4098$). Therefore, y^6 is a very small positive number, less than 1. Therefore, $x^6 = 13536 - y^6$. Since y^6 is positive, $x^6 < 13536$. Thus, the integer part of x^6 is 13535. Therefore, the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$ is 13535. **Final Answer**: 13535 </think> To find the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$, we denote $x = \sqrt{7} + \sqrt{5}$ and $y = \sqrt{7} - \sqrt{5}$. We know that $x \cdot y = 2$, $x + y = 2\sqrt{7}$, and $x - y = 2\sqrt{5}$. We use the recurrence relation for $x^n + y^n$: $x^1 + y^1 = 2\sqrt{7}$, $x^2 + y^2 = 24$, $x^3 + y^3 = 44\sqrt{7}$, $x^4 + y^4 = 568$, $x^5 + y^5 = 1048\sqrt{7}$, $x^6 + y^6 = 13536$. Since $y = \sqrt{7} - \sqrt{5}$ is less than 1, y^6 is a very small positive number. Therefore, $x^6 = 13536 - y^6$, which is slightly less than 13536. Thus, the greatest integer less than $(\sqrt{7} + \sqrt{5})^6$ is **13535**.

Final Answer: 13535.

Figure 12: Case study under the DeepSeek-R1-Distill-Qwen-1.5B model with three settings: Base, PPO, and Bingo. Blue highlights some redundant and repetitive tokens, while red marks omitted content and the final answer.

Problem: In how many ways can 8 people sit around a round table if 3 of the people – Pierre, Rosa, and Thomas – all want to sit together? (Two seatings are considered the same if one is a rotation of the other.)

✗ Response of Base Model with 5783 words:

Okay, so I have this function $f(x) = \sqrt{(x-1)}/(x-2)$, and I need to find the smallest possible integer value for x such that $f(x)$ is a real number. Hmm, let me think about how to approach this. First, I remember that for a square root to be real, the expression inside the square root, which is called the radicand, must be greater than or equal to zero. So, in this case, the radicand is $(x-1)$. That means $x-1$ has to be greater than or equal to zero. Let me write that down: $x-1 \geq 0$ Solving for x , I add 1 to both sides: $x \geq 1$ Okay, so x has to be at least 1. But wait, there's also the denominator in the function, which is $(x-2)$. For the function to be defined, the denominator can't be zero because division by zero is undefined. So, I need to make sure that $x-2$ is not equal to zero. Let me write that condition: ... But x can't be 2, so the next integer after 1 is 2, which is excluded. So, the next integer is 3. wait, is 3 the smallest integer that satisfies both conditions? Let me check. ... wait, if x is 1, then $f(x)$ is 0, which is real. So, 1 is a valid value. But is 1 the smallest integer? ... wait, but hold on, the problem says "the smallest possible integer value for x ." So, 1 is the smallest integer in the domain. So, 1 is the answer. But wait, hold on, I just thought of something. The function $f(x)$ is defined as $\sqrt{(x-1)}/(x-2)$. So, for $x=1$, it's 0, which is fine. But is 1 the smallest integer? Or is there a smaller integer? But integers go down to negative infinity, but x can't be less than 1 because of the square root. So, 1 is the smallest possible integer. ... So, 1 is the smallest integer in the domain, so that's the answer. But just to make sure, let me think about the function's domain. ... But wait, hold on, I just thought of something else. The problem is asking for the smallest possible integer value for x such that $f(x)$ has a real number value. So, 1 is the smallest integer, but is 1 the smallest integer in the domain? Yes, because x can't be less than 1. But wait, let me think about the function again. If x is 1, $f(x)$ is 0, which is real. So, 1 is acceptable. So, 1 is the answer. ... So, the integers in the domain are 1, 3, 4, 5, etc. So, the smallest integer is 1. Therefore, the [Generation Termination without final answer]

✓ Response of PPO with 437 words:

To determine the smallest possible integer value for x such that $f(x) = \frac{\sqrt{x-1}}{x-2}$ has a real number value, we need to consider the conditions under which the function is defined. 1. **Numerator Condition:** The expression inside the square root, $x-1$, must be non-negative because the square root of a negative number is not a real number. $x-1 \geq 0 \Rightarrow x \geq 1$ 2. **Denominator Condition:** The denominator, $x-2$, must not be zero because division by zero is undefined. $x-2 \neq 0 \Rightarrow x \neq 2$ 3. **Combined Domain:** Combining the above conditions, x must be greater than or equal to 1 but not equal to 2. $x \geq 1$ and $x \neq 2$ 4. **Smallest Integer Value:** The smallest integer greater than or equal to 1 is 1. However, we need to check if $x=1$ satisfies the conditions. - For $x=1$: $f(1) = \frac{\sqrt{1-1}}{1-2} = \frac{0}{-1} = 0$ Since 0 is a real number, $x=1$ is valid. Therefore, the smallest possible integer value for x is 1. To determine the smallest possible integer value for x such that $f(x) = \frac{\sqrt{x-1}}{x-2}$ has a real number value, let's analyze the function step by step. 1. **Numerator Condition:** - The expression inside the square root, $x-1$, must be non-negative because the square root of a negative number is not a real number. $x-1 \geq 0 \Rightarrow x \geq 1$ 2. **Denominator Condition:** - The denominator, $x-2$, must not be zero because division by zero is undefined. $x-2 \neq 0 \Rightarrow x \neq 2$ 3. **Combined Domain:** - Combining the above conditions, x must be greater than or equal to 1 but not equal to 2. $x \geq 1$ and $x \neq 2$ 4. **Smallest Integer Value:** - The smallest integer greater than or equal to 1 is 1. However, we need to check if $x=1$ satisfies the conditions. - For $x=1$: $f(1) = \frac{\sqrt{1-1}}{1-2} = \frac{0}{-1} = 0$ Since 0 is a real number, $x=1$ is valid. Therefore, the smallest possible integer value for x is: 1

✓ Response of Bingo with 155 words:

First, I need to determine the domain of the function $f(x) = \frac{\sqrt{x-1}}{x-2}$ to find the smallest integer x for which $f(x)$ is a real number. I know that for the square root to be real, the expression inside must be non-negative. So, $x-1 \geq 0$, which implies $x \geq 1$. Additionally, the denominator $x-2$ cannot be zero, so $x \neq 2$. Combining these conditions, the domain is $x \geq 1$ with $x \neq 2$. The smallest integer greater than or equal to 1 is 1, but I need to check if $x=1$ is valid. Plugging in $x=1$, the denominator becomes -1 , which is allowed, and the square root is 0. So $f(1) = 0$, which is a real number. Therefore, the smallest possible integer value for x is 1. </think> The function $f(x) = \frac{\sqrt{x-1}}{x-2}$ is defined when the expression inside the square root is non-negative and the denominator is not zero. 1. ****Square Root Condition:**** $x-1 \geq 0$. $x \geq 1$. 2. ****Denominator Condition:**** $x-2 \neq 0$. $x \neq 2$. Combining these conditions, the domain of $f(x)$ is $x \geq 1$ with $x \neq 2$. The smallest integer value for x that satisfies these conditions is $x=1$.

Final Answer: 1.

Figure 13: Case study under the DeepSeek-R1-Distill-Qwen-7B model with three settings: Base, PPO, and Bingo. Blue highlights some redundant and repetitive tokens, while red marks omitted content and the final answer.

Table 9: Notation used throughout the paper

Notation	Description
<i>General</i>	
y	Sequence of tokens generated by the language model
x	Input prompt for the language model
n	Total length of the sequence y
y_i	i -th token in the generated sequence y
$\hat{z}(y)$	Extracted final answer from the generated sequence y
z	Ground-truth answer
\mathbb{E}_{π_θ}	Expectation over policy π_θ
$A(L)$	Expected accuracy as a function of output length L
L	Length of the output sequence generated by the model
L_{\max}	Maximum response length in the dataset
Acc	Exact match accuracy of the final output
L-Acc	Length-normalized accuracy, defined as $\text{Acc} \times \sqrt{1 - \frac{L}{L_{\max}}}$
$S(y_i)$	Significance score of token y_i
L^s	Number of significant tokens in the response
L^{is}	Number of insignificant tokens in the response
τ	Threshold for classifying a token as significant or insignificant
<i>Reinforcement Learning</i>	
π_θ	Policy parameterized by θ
\hat{A}_t	Advantage estimate at time step t
$r_t(\theta)$	Importance sampling ratio for policy optimization
R^{BINGO}	Reward function in the BINGO framework
$\mathcal{J}_{\text{BINGO}}(\theta)$	PPO objective with BINGO reward function
ϵ	Clipping parameter in the PPO objective
<i>Rewards and Penalties</i>	
$r_{is}(y)$	Reward for insignificant tokens in sequence y
$r_s(y)$	Reward for significant tokens in sequence y
λ_c	Coefficient for penalty on correct responses
λ_w	Coefficient for penalty on incorrect responses
k	Dynamic scaling factor for length reward
α	Scaling factor for the decay in dynamic length reward
β	Threshold for transition between exploration and compression in dynamic reward
<i>Length Penalty</i>	
L_{ref}^{is}	Reference number of insignificant tokens
L_{ref}^s	Reference number of significant tokens
$k(t)$	Dynamic scaling factor for adjusting length reward over time
<i>Miscellaneous</i>	
\mathcal{M}_e	Model used to estimate token significance (LLMLingua-2)
$\mathbb{1}[\cdot]$	Indicator function (1 if true, 0 otherwise)
\mathcal{Y}_{sig}	Set of significant tokens in the sequence y
$\mathcal{Y}_{\text{insig}}$	Set of insignificant tokens in the sequence y