

Stable Adaptive Thinking via Advantage Shaping and Length-Aware Gradient Regulation

Zihang Xu^{1,2*}, Haozhi Xie^{1*}, Ziqi Miao^{1,2*}, Chen Qian^{2,3}, Lijun Li^{2†}

¹Beihang University, ²Shanghai Artificial Intelligence Laboratory,
³Gaoling School of Artificial Intelligence, Renmin University of China

lilijun@pjlab.org.cn

Abstract

Large reasoning models (LRMs) achieve strong performance through extended reasoning traces, but they often exhibit overthinking behavior for low-complexity queries. Existing efforts to mitigate this issue are fundamentally limited by unstable accuracy–efficiency trade-offs and poor robustness to heterogeneous reasoning behaviors. To address these challenges, we propose a two-stage framework for stable adaptive thinking in LRMs. The framework first applies Hybrid Fine-Tuning to expose the model to both thinking and no-thinking behaviors, establishing well-conditioned initialization. It then performs adaptive reinforcement learning with Correctness-Preserving Advantage Shaping (CPAS) to avoid suppressing correct long-chain reasoning, and Length-Aware Gradient Regulation (LAGR) to stabilize optimization under severe reasoning-length heterogeneity. Extensive experiments on Qwen2.5-1.5B and 7B show consistent improvements over strong baselines, achieving up to +3.7/+3.6 accuracy points while reducing generated tokens by 40.6%/43.9%. Further analyses across varying problem difficulties and out-of-distribution tasks confirm the robustness and generalization of our approach.

1 Introduction

Recent advancements in large reasoning models (LRMs), such as OpenAI-o1 (OpenAI et al., 2024) and DeepSeek-R1 (Guo et al., 2025), have demonstrated remarkable progress in solving complex tasks. By explicitly generating extended reasoning traces, these models exhibit strong reasoning and generalization capabilities across a wide range of domains (Xu et al., 2025). Although long-form reasoning benefits challenging problems, existing LRMs lack difficulty-aware control ability and thus often produce unnecessarily long reasoning chains even for simple queries (Yue et al., 2025; Liu et al.,




Figure 1: Comparison of efficient reasoning strategies.

2025). This overthinking issue significantly increases token usage and inference costs, posing a critical bottleneck for the practical efficiency of LRMs (Chen et al., 2025b).

Existing efficient reasoning methods reduce overthinking through length control, routing, or adaptive policy learning (Feng et al., 2025; Sui et al., 2025), as illustrated in Figure 1. Length-based penalties have been introduced to discourage excessively long reasoning chains (Luo et al., 2025a; Aggarwal and Welleck, 2025; Xia et al., 2025; Yi et al., 2025; Deng et al., 2025); however, these approaches operate within the thinking mode, and the additional constraints in these approaches may limit the exploration space of LRMs. Routing-based methods reduce inference cost by selecting between reasoning and non-reasoning models (Hu et al., 2024; Ong et al., 2025; Singh et al., 2026), but they typically depend on external annotations for supervision and the maintenance of multiple models, thereby increasing both training and deployment cost. More recently, reinforcement-learning-based adaptive thinking methods aim to

*Equal contribution. †Corresponding author.

improve efficiency by encouraging shorter reasoning chains. ThinkLess (Fang et al., 2025) achieves this via a decoupled GRPO formulation, whereas AdapThink (Zhang et al., 2025a) employs a PPO-based objective to promote short-chain reasoning. Despite their effectiveness, these methods often face difficulties in achieving a stable balance between reasoning accuracy and efficiency during training.

These observations highlight two key challenges in adaptive thinking. First, there is an inherent trade-off between efficiency and accuracy. While adaptive thinking aims to reduce computational costs, an excessive bias toward brevity may constrain the model’s capacity for deep reasoning on complex tasks, ultimately degrading performance. Second, severe length heterogeneity during training poses additional difficulties. Within a single rollout group, reasoning chains may range from a few hundred tokens to tens of thousands of tokens, resulting in uneven optimization granularity and imbalanced gradient contributions, which may compromise training stability.

To address these challenges, we propose Stable Adaptive Thinking via Advantage Shaping and Length-Aware Gradient Regulation, a two-stage framework for stable and effective adaptive reasoning in large reasoning models. The proposed approach enables the model to adjust its reasoning depth according to input difficulty, improving efficiency on simple queries while preserving accuracy on complex ones. In the first stage, **Hybrid Fine-Tuning** is employed to establish well-conditioned initialization by exposing the model to both short and long reasoning behaviors. In the second stage, adaptive thinking is optimized through reinforcement learning with two key mechanisms. **Correctness-Preserving Advantage Shaping**, in contrast to prior approaches that directly reward short chains, introduces an advantage shaping strategy that avoids suppressing correct long-chain reasoning trajectories when encouraging efficiency, thereby preserving the model’s exploration capacity. **Length-Aware Gradient Regulation** addresses optimization instability arising from the severe reasoning length heterogeneity by regulating gradient allocation across different responses, while ensuring that control signals remain effective under long and heterogeneous reasoning sequences, leading to more stable training dynamics. These components provide a principled solution for stable adaptive thinking under efficiency–accuracy trade-

offs. In general, the contributions of this paper can be summarized as follows:

- We address key optimization challenges in adaptive thinking. Specifically, we focus on balancing computational efficiency with reasoning accuracy, and on mitigating training instability arising from severe length heterogeneity across different reasoning modes during reinforcement learning optimization.
- We propose a two-stage framework for stable adaptive thinking. The framework integrates hybrid fine-tuning with adaptive reinforcement learning, while employing correctness-preserving advantage shaping and length-aware gradient regulation for stable training.
- Extensive experiments on the Qwen2.5 series across multiple model scales and mathematical benchmarks demonstrate that our approach consistently outperforms existing baselines, achieving up to +3.7/+3.6 accuracy improvements while reducing the number of generated tokens by 40.6%/43.9%. Further analyses across varying problem difficulties and diverse tasks confirm the robustness and generalization of our approach.

2 Preliminaries

2.1 Adaptive Thinking

Adaptive thinking allows a single policy to decide whether explicit reasoning is needed before producing an answer. Given an input problem x , the model produces an output y consisting of a chain-of-thought trajectory S and a final answer a : $y = \{S, a\}$. We introduce a control token to specify the thinking mode $m \in \mathcal{M} = \{\text{/think}, \text{/no_think}\}$: when the control token is /think , the model generates a non-empty chain-of-thought trajectory S in the thinking mode; when the control token is /no_think , the chain-of-thought is empty $S = \emptyset$ and the model directly produces a .

We implement adaptive thinking with a single policy model. Given input $x \sim \mathcal{D}$, the model first generates a control token $m \in \mathcal{M}$, and then generates the remaining output conditioned on m . Thus, the same parameters θ define both mode selection π_θ (the distribution of the first control token) and response generation $P_\theta(y \mid x, m)$.




Figure 2: Overview of our two-stage training pipeline. Stage 1 performs Hybrid Fine-Tuning (HFT) on paired thinking and no-thinking formats to initialize a unified policy. Stage 2 applies GRPO-style reinforcement learning with correctness-preserving advantage shaping (CPAS) and length-aware gradient regulation (LAGR) to stabilize optimization under extreme length heterogeneity and to learn when to think.

2.2 Problem Statement

The goal of adaptive thinking is to reduce unnecessary generation while preserving correctness. We define a utility function that trades off accuracy and efficiency. For an output $y = \{S, a\}$, we use an accuracy reward $R_{\text{acc}}(y | x) \in \{0, 1\}$ indicating whether the final answer a is correct (verified by an automatic checker), and a length reward $R_{\text{len}}(y | x)$ that favors shorter outputs. The overall utility is

$$U(x, y) = R_{\text{acc}}(y | x) + \gamma R_{\text{len}}(y | x)$$

where $\gamma > 0$ controls the efficiency–accuracy trade-off. Our objective is to learn a policy that selects m and generates y to maximize expected utility:

$$\max_{\theta} E_{x \sim \mathcal{D}} \left[E_{m \sim \pi_{\theta}(\cdot | x)} E_{y \sim P_{\theta}(\cdot | x, m)} [U(x, y)] \right].$$

3 Methodologies

3.1 Overall Framework

We propose a novel framework for *stable adaptive thinking*, which enables a single model to dynamically regulate its thinking depth according to problem difficulty. The objective is to avoid unnecessary overthinking on simple queries while preserving strong deliberative capability on challenging problems. Starting from a pretrained base language model, our framework adopts a two-stage training paradigm that comprises hybrid fine-tuning

for capability warm-up and reinforcement learning for stable adaptive thinking. An overview of the training pipeline is illustrated in Figure 2.

Stage I: Hybrid Fine-Tuning. We first fine-tune the base model on large-scale mathematical reasoning data with heterogeneous supervision, including long CoT solutions and short direct answers distilled from strong teacher models. This stage injects both thinking and no-thinking capabilities into a unified model and provides a stable warm-up for subsequent reinforcement learning.

Stage II: Stable Adaptive Thinking with Reinforcement Learning. We then optimize the model with a modified GRPO algorithm, allowing it to explore both thinking and no-thinking trajectories and learn *when to think*. To address length heterogeneity and prevent mode collapse, we introduce correctness-preserving advantage shaping and length-aware gradient regulation, which stabilize exploration and optimization.

3.2 Hybrid Fine-Tuning for Warm-Up

Hybrid Fine-Tuning (HFT) provides a stable cold start for adaptive reinforcement learning by exposing the model to thinking and no-thinking responses in a unified format. Rather than learning the adaptive policy, HFT broadens the solution space to cover heterogeneous response behaviors.

Hybrid-Formatted Data Construction. We construct the HFT dataset from large-scale open-source mathematical corpora by distilling two response modes. Thinking-mode examples are generated by DeepSeek-R1-0528 with the `/think` prefix and a non-empty `<think>...</think>` block, while no-thinking-mode examples are generated by DeepSeek-V3-0324 with the `/no_think` prefix and an empty `<think></think>` block. This hybrid format provides explicit supervision for both deliberative reasoning and direct answering. More details are provided in Appendix A.1.

Optimization Objective. During HFT, the model is optimized using standard supervised learning objectives over both response types, without imposing any explicit preference between thinking and no-thinking outputs. Given the hybrid-formatted dataset $\mathcal{D}_{\text{HFT}} = \{(x_i, y_i)\}_{i=1}^N$, the model is optimized to maximize the conditional likelihood of the target response given the input:

$$\mathcal{L}_{\text{HFT}} = -E_{(x,y) \sim \mathcal{D}_{\text{HFT}}} \left[\sum_{t=1}^{|y|} \log p_{\theta}(y_t \mid x, y_{<t}) \right].$$

Role of HFT. HFT makes both thinking and no-thinking trajectories accessible to the policy, defining the feasible action space for reinforcement learning. This warm-up reduces exploration difficulty and yields a better-conditioned initialization.

3.3 Stable Adaptive thinking with Reinforcement Learning

After HFT, the model is further trained to acquire *adaptive thinking* capability through reinforcement learning. While HFT equips the model with both thinking-mode and no-thinking-mode generation behaviors, it does not enable the model to autonomously decide which mode to adopt for a given query. This decision-making ability is learned in the reinforcement learning stage.

Specifically, we build upon GRPO and train the model to explore both thinking and no-thinking trajectories for each query. In this setting, the policy learns to decide *whether to engage in explicit thinking* by selecting between thinking-mode and no-thinking-mode responses. However, directly applying vanilla GRPO often leads to training instability and mode collapse. We identify two key challenges that arise from GRPO under heterogeneous thinking trajectories, and introduce corresponding mechanisms to address them.

Correctness-Preserving Advantage Shaping.

Existing adaptive thinking approaches often discourage overthinking by assigning higher rewards to shorter thinking trajectories. While this efficiency preference can reduce token usage, it interacts nontrivially with GRPO-style group-wise normalization. Specifically, advantages are computed by normalizing rewards within a rollout group sampled for the same input: $A_i = \frac{r_i - \text{Mean}(r)}{\text{Std}(r)}$, where r_i denotes the scalar reward used by the adaptive thinking objective. When shorter trajectories systematically receive higher rewards (Aggarwal and Welleck, 2025; Fang et al., 2025), this normalization induces a relative comparison bias: longer but correct thinking trajectories may satisfy $0 < r_i < \text{Mean}(r)$ and thus receive $A_i < 0$. Consequently, valid thinking behaviors are actively suppressed during optimization, limiting exploration and potentially leading to mode collapse toward no-thinking responses.

To alleviate this issue, we introduce *Correctness-Preserving Advantage Shaping* (CPAS), which reshapes advantages for correct responses based on their thinking mode. Let A_i^{GRPO} denote the advantage computed by GRPO for response o_i , and let $r_i \in \{0, 1\}$ denote the reward. The final advantage is defined as

$$A_i = \begin{cases} A_i^{\text{GRPO}} + \delta, & \text{if } r_i = 1 \text{ and } o_i \text{ is short-chain,} \\ A_i^{\text{GRPO}}, & \text{otherwise.} \end{cases}$$

where, $\delta > 0$ is a small reward coefficient that provides additional encouragement for correct short-chain responses, while leaving correct long-chain responses unaffected.

Role of CPAS. By reshaping advantages for correct responses, CPAS ensures that correct thinking trajectories receive non-negative advantage signals. This prevents correct long-chain responses from being inadvertently penalized, thereby preserving the model’s exploration capacity. Consequently, CPAS contributes to more stable training dynamics and enhances the effectiveness of adaptive thinking.

Length-Aware Gradient Regulation. Under GRPO, policy optimization is performed by sampling a group G of responses $\{o_i\}_{i=1}^G$ for each query q and aggregating token-level policy gradients within each response. The response-level

gradient contribution of o_i can be written as

$$g_i^{\text{GRPO}} = \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \hat{A}_{i,t} \nabla_{\theta} \log \pi_{\theta}(o_{i,t} \mid q, o_{i,<t}),$$

where $|o_i|$ denotes the response length and $\hat{A}_{i,t}$ is the token-level advantage.

In adaptive thinking, rollout groups often exhibit severe length heterogeneity, ranging from short no-thinking responses to long thinking trajectories. The $1/|o_i|$ normalization systematically attenuates gradient magnitudes for long responses, increases gradient variance, and biases optimization toward short trajectories, which may lead to premature collapse of thinking behaviors.

To mitigate this imbalance, we propose *Length-Aware Gradient Regulation*, which explicitly reweights response-level gradient contributions according to response length. Specifically, LAGR employs a length-dependent weighting scheme:

$$g_i^{\text{LAGR}} = w_i^{(\beta)} \sum_{t=1}^{|o_i|} \hat{A}_{i,t} \nabla_{\theta} \log \pi_{\theta}(o_{i,t} \mid q, o_{i,<t}),$$

$$w_i^{(\beta)} = \frac{1}{M} \cdot \frac{|o_i|^{-\beta}}{\sum_{j=1}^G |o_j|^{-\beta}}, \quad \beta \in [0, 1].$$

where M is a constant scaling factor used to control the overall magnitude of the gradient. When $\beta = 1$, LAGR recovers the GRPO-like length-sensitive behavior that favors short responses, while $\beta = 0$ assigns uniform weights across responses, thereby alleviating the suppression of long ones. By adjusting β , LAGR provides a mechanism to balance optimization stability and thinking flexibility.

While LAGR mitigates length-induced imbalance at the response level, adaptive thinking introduces an additional challenge at the token level. In practice, the decision of whether to engage in thinking is primarily encoded by the prefix control token. However, when responses are long, the gradient signal of this control token can be diluted by the large number of subsequent tokens. To preserve this adaptive control signal, we further apply an explicit boosting to the control token. Let $t = 0$ denote the control token. The final response-level gradient under LAGR is given by

$$\nabla_{\theta} \mathcal{J}(\theta) = \nabla_{\theta} \sum_{i=1}^G w_i^{(\beta)} \left[\lambda \ell_{i,0}^{\text{clip}} + \sum_{t=1}^{|o_i|} \ell_{i,t}^{\text{clip}} \right],$$

$$\ell_{i,t}^{\text{clip}} = \min \left(r_{i,t}(\theta) \hat{A}_{i,t}, \right. \\ \left. \text{clip}(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t} \right).$$

where λ is a control-token boosting factor that compensates for gradient dilution in long responses.

Role of LAGR. LAGR rebalances policy gradients under length heterogeneity, mitigating GRPO’s short-trajectory bias. By preserving effective gradient signals in long trajectories and boosting the control token, LAGR prevents decision-signal dilution and stabilizes adaptive thinking.

4 Experiments

4.1 Experimental Settings

Models. We adopt Qwen-2.5-1.5B-Base and Qwen-2.5-7B-Base as the initial policy models (Qwen et al., 2025). Both are extensively pre-trained and serve as the foundation for jointly learning thinking and no-thinking behaviors.

Datasets and Metrics. For HFT, we build a large-scale corpus from open-source mathematical datasets, including OpenR1-Math-220k (Lozhkov et al., 2025), DeepMath-103K (He et al., 2025b), and NuminaMath (LI et al., 2024), and obtain paired supervision by distilling thinking-mode responses from DeepSeek-R1-0528 (Guo et al., 2025) and no-thinking responses from DeepSeek-V3-0324 (DeepSeek-AI et al., 2025). For RL, we primarily adopt the DeepScaleR (Luo et al., 2025b) dataset. We evaluate on MATH-500 (Lightman et al., 2024), AIME-2024¹, and AIME-2025², and further assess OOD generalization on GPQA (Rein et al., 2024) and MMLU (Hendrycks et al., 2021). We report accuracy, average token length, and no-thinking ratio to jointly measure model performance and inference efficiency. Additional results on GSM8K (Cobbe et al., 2021) and AMC (Yang et al., 2024) are provided in Appendix B.

Baselines. To ensure a comprehensive comparison, we evaluate our method at both 1.5B and 7B scales against four categories of baselines. (1) *Base Variants*: thinking-only, no-thinking-only, and HFT hybrid models derived from the same base architectures. (2) *Length-Penalization Methods*: O1-Pruner (Luo et al., 2025a), ShorterBetter (Yi et al., 2025), and S-GRPO (Dai et al., 2025), which improve efficiency by penalizing overly long reasoning trajectories within the thinking mode. (3) *Routing Methods*: RouteLLM (Ong et al., 2025), which

¹<https://huggingface.co/datasets/AI-MO/aimo-validation-aimo>

²<https://huggingface.co/datasets/math-ai/aimo25>

Table 1: Overall performance on MATH-500, AIME-2024, and AIME-2025 for Qwen2.5 series models. Len is the average number of generated tokens per example. $Ratio_{NT}$ denotes the fraction of examples solved in NoThinking mode. For methods without explicit NoThinking routing, $Ratio_{NT}$ is set to 0.0%. ΔAcc and ΔLen are computed relative to the always-Thinking baseline of the same base model.

Method	MATH-500			AIME-2024			AIME-2025			Average	
	Acc	Len	$Ratio_{NT}$	Acc	Len	$Ratio_{NT}$	Acc	Len	$Ratio_{NT}$	ΔAcc	ΔLen
Qwen-2.5-1.5B											
Thinking	80.4	10441	0.0%	25.7	35829	0.0%	20.6	36977	0.0%	0.0	+0.0%
No-Thinking	60.4	2320	100.0%	9.1	10751	100.0%	7.5	11150	100.0%	-16.6	-70.9%
HFT	72.3	7106	42.4%	16.8	23673	54.2%	14.5	25376	46.2%	-7.7	-32.5%
O1-Pruner	77.1	7043	0.0%	20.3	33417	0.0%	16.7	32185	0.0%	-4.2	-12.7%
ShorterBetter	77.9	4605	0.0%	22.1	24107	0.0%	18.6	<u>22850</u>	0.0%	-2.7	-38.1%
S-GRPO	79.7	6638	0.0%	25.2	27849	0.0%	20.1	29653	0.0%	-0.6	-23.0%
RouteLLM	76.7	3862	81.6%	19.8	28305	30.8%	17.4	29973	27.8%	-4.3	-25.4%
AdaptThink	<u>81.1</u>	4341	69.4%	<u>25.6</u>	22819	4.2%	<u>22.7</u>	24187	3.3%	+0.9	-38.3%
Thinkless	78.5	4473	71.3%	24.5	25294	2.9%	19.3	26351	2.5%	-1.5	-32.6%
Ours	84.6	<u>4176</u>	69.8%	28.1	<u>23258</u>	3.7%	25.2	22062	5.4%	+3.7	-40.6%
Qwen-2.5-7B											
Thinking	90.6	8309	0.0%	65.6	29936	0.0%	59.3	35217	0.0%	0.0	+0.0%
No-Thinking	75.1	1890	100.0%	25.5	10464	100.0%	20.2	9921	100.0%	-31.6	-69.7%
HFT	86.3	5619	45.2%	50.9	23484	42.5%	44.2	25017	40.4%	-11.3	-26.3%
O1-Pruner	87.9	5567	0.0%	62.4	25977	0.0%	57.5	32451	0.0%	-2.6	-12.9%
ShorterBetter	88.1	3594	0.0%	61.9	18615	0.0%	57.1	21733	0.0%	-2.8	-40.2%
S-GRPO	90.9	5326	0.0%	65.1	22893	0.0%	59.7	28474	0.0%	+0.1	-22.8%
RouteLLM	88.2	3121	81.8%	55.3	24011	31.2%	51.7	28387	27.1%	-6.8	-24.4%
AdaptThink	<u>91.1</u>	3462	77.8%	<u>65.2</u>	<u>18123</u>	12.1%	<u>59.1</u>	22544	10.4%	0.0	-39.9%
Thinkless	89.8	<u>3302</u>	78.6%	64.0	18172	13.8%	60.1	<u>21463</u>	12.5%	-0.5	-41.6%
Ours	92.4	3403	77.6%	69.2	17762	12.5%	64.6	20060	9.6%	+3.6	-43.9%

trains an external router to choose between separate thinking and no-thinking models. (4) *Adaptive Thinking Methods*: ThinkLess (Fang et al., 2025) and AdaptThink (Zhang et al., 2025a), which learn hybrid thinking policies via GRPO- or PPO-based optimization and are initialized from HFT models. All baselines use consistent base architectures and appropriate training settings for fair comparison.

Implementation Details. All experiments are conducted on 16 NVIDIA H100 GPUs. In the Hybrid Fine-Tuning stage, we set the maximum context length to 16K tokens, with overlong samples truncated accordingly. Models are trained for 3 epochs with AdamW, using a 10% linear warmup followed by cosine decay and a peak learning rate of 1×10^{-4} . For Reinforcement Learning, we use VeRL (Sheng et al., 2025) with a 24K context length and optimize the policy model using AdamW with a constant learning rate of 1×10^{-6} . We set the batch size to 128, the micro-batch size to 64, and sample 8 responses per query during rollout. Correctness feedback is provided by Math-Verify. More details are in Appendix A.3.

4.2 Main Results

In this section, we conduct a comprehensive evaluation of different models on three mathematical reasoning benchmarks. As shown in Table 1, our method achieves consistent gains at both the 1.5B and 7B scales, improving performance by 3.7 and 3.6 points, respectively, while reducing token costs by 40.6% and 43.9%, demonstrating both accuracy and efficiency improvements. Compared to length-penalty methods, our approach leverages an explicit no-thinking mode that does not constrain exploration, thereby achieving lower token consumption and improved overall performance. Compared to external routing methods, our approach requires no additional annotated data or deployment overhead, and can more accurately select appropriate thinking modes based on the model’s capability. Finally, relative to existing adaptive thinking methods, the proposed correctness-preserving advantage shaping and length-aware gradient regulation achieve significantly better performance under comparable generation lengths, enabling stable and effective adaptive thinking. Additional results on




Figure 3: Difficulty-aware mode selection and performance. (a) Performance of Two Think Mode on MATH-500, AIME-2024, and AIME-2025. (b) Mode ratio across MATH-500 difficulty levels. (c) Accuracy across difficulty levels, comparing our adaptive policy with always-Thinking and always-No-Thinking baselines.




Figure 4: Ablation and sensitivity analysis. (a) Training dynamics with/without CPAS: mean response length (left) and AIME-2024 accuracy (right). (b) Effect of the LAGR length-weight parameter β (left) and the control-token boost factor λ (right) on accuracy and the no-thinking ratio.

broader mathematical reasoning datasets and alternative backbone architectures are provided in Appendix B, further supporting our robustness.

4.3 Analysis of Adaptive Thinking

We further analyze model behavior across problems of varying difficulty to evaluate the effectiveness of adaptive thinking. As shown in Figure 3(a), on the relatively easy MATH-500 benchmark, the proposed method selects the no-thinking mode for up to 77.6% of the queries, whereas on the more challenging AIME-2024 and AIME-2025 benchmarks, it activates the thinking mode more frequently. Notably, performance under the no-thinking mode consistently surpasses that under the thinking mode, indicating that the model learns to adaptively select the more appropriate strategy according to problem difficulty. Figures 3(b) and 3(c) further illustrate the model’s thinking behavior across different difficulty levels within MATH-500. As the difficulty increases, the model progressively increases its reliance on the thinking mode. Compared with the original model that exclusively uses either the thinking or no-thinking mode, our method consistently achieves superior performance across all difficulty

levels. These results demonstrate that the proposed approach attains a more favorable balance between reasoning efficiency and accuracy.

4.4 Ablation and Sensitivity Analysis

Effect of CPAS. Figure 4(a) illustrates the training dynamics with and without CPAS. Compared with the baseline, our method exhibits a faster increase in response length at early stages and reaches a higher peak, indicating earlier and deeper exploration. The length then gradually decreases and converges to a similar level, achieving stable adaptive thinking behavior. Evaluation on the AIME benchmarks further shows that our method consistently outperforms the baseline, demonstrating a better balance between accuracy and efficiency, which confirms the effectiveness of CPAS.

Effect of LAGR. We analyze the impact of LAGR under different parameters. As shown in Figure 4(b) left, values of β closer to 1 encourage no-thinking chains, whereas values closer to 0 promote the thinking exploration. An intermediate setting of $\beta = 0.4$ provides the best trade-off between accuracy and efficiency. Figure 4(b) right examines the effect of control-token weighting. When

Table 2: Out-of-distribution generalization results on GPQA and MMLU datasets.

Method	Acc	Len	Ratio _{NT}	Δ Len
GPQA				
Thinking	47.5	24135	0.0%	–
No-Thinking	35.3	8031	100.0%	-66.7%
HFT	41.8	16849	44.1%	-30.2%
Ours	50.4	11826	31.3%	-51.0%
MMLU				
Thinking	67.6	3769	0.0%	–
No-Thinking	52.2	1137	100.0%	-69.8%
HFT	59.7	2263	53.1%	-40.0%
Ours	71.4	1624	41.3%	-56.9%

the weight is set to 1, the ratio between thinking and no-thinking modes remains largely unchanged, indicating that the control signal is diluted. Increasing the weight to 10 enables stable adaptive thinking, while larger weights induce overly aggressive updates, prematurely assigning solvable samples to the thinking mode and reducing efficiency.

4.5 Generalization Analysis

We evaluate the generalization ability of our method on two challenging benchmarks: GPQA and MMLU. GPQA consists of graduate-level multiple-choice questions across diverse scientific domains, while MMLU further covers 57 subjects spanning STEM, humanities, and social sciences, thereby providing a broader testbed for assessing cross-domain robustness. As shown in Table 2, our method consistently achieves the best performance on both benchmarks. On GPQA, it reduces the average response length by 51.0%, with 31.3% of queries handled in the no-thinking mode, while maintaining strong performance. On MMLU, our method further improves accuracy from 67.6 to 71.4 and reduces the average response length by 56.9%, activating no-thinking mode for 41.3% of the queries. These results indicate that our approach can generalize beyond the training distribution and maintain a favorable accuracy–efficiency trade-off across both specialized scientific reasoning and broader knowledge-intensive tasks.

5 Related Work

5.1 Large Reasoning Models

Large reasoning models (OpenAI et al., 2024; Guo et al., 2025) have recently attracted substantial attention due to their superior performance. Enabled by scalable reinforcement learning (He et al., 2025b; Yu et al., 2025a; Zheng et al., 2025; Miao

et al., 2026; Wang et al., 2026; Feng et al., 2026), these models can generate long reasoning chains using only outcome-based and format rewards, leading to improved performance in complex reasoning tasks. To further generalize in various scenarios, previous work has explored reinforcement learning with verifiable rewards (Yu et al., 2025b) and introduced explicit supervision over the reasoning process (Zhang et al., 2025b). However, most existing approaches overlook the efficiency issues caused by long redundant reasoning chains, which increases the inference cost. Creating efficient reasoning without sacrificing performance remains an open challenge.

5.2 Efficient Reasoning for LRMs

Efficient reasoning for LRMs aims to reduce overthinking while preserving accuracy. One line of work focuses on compressing or regularizing reasoning traces via prompting, supervised tuning with variable-length CoT, or length-regularized RL (Aytes et al., 2025; Chen et al., 2025b; Ma et al., 2025; Aggarwal and Welleck, 2025; Luo et al., 2025a). These methods typically assume reasoning is always necessary and may restrict the exploration. Another line studies hybrid thinking via routing, allocating compute by switching between thinking and no-thinking models based on task difficulty (He et al., 2025a; Liang et al., 2025; Singh et al., 2026), but it often requires external supervision and deployment costs. More recently, RL-based adaptive thinking learns when to think using explicit control tokens or rewards for no-thinking behaviors (Fang et al., 2025; Jiang et al., 2025; Zhang et al., 2025a; Lou et al., 2025; Chen et al., 2025a). However, achieving stable optimization under the efficiency–accuracy trade-off remains challenging, motivating our stability-oriented approach.

6 Conclusion

In this work, we address the training instability caused by heterogeneous reasoning lengths in existing adaptive thinking methods, as well as the difficulty of jointly optimizing reasoning accuracy and efficiency. We propose a novel two-stage framework for stable adaptive thinking. In Stage I, we adopt HFT to endow the model with both thinking and no-thinking modes. In Stage II, CPAS preserves the model’s ability to explore long reasoning chains, while LAGR stabilizes reinforcement learn-

ing training. Extensive experiments demonstrate that our approach achieves a better balance between performance and efficiency, underscoring the potential of our proposed advantage-based model optimization strategy for reinforcement learning.

Limitations

In this section, we discuss several limitations of this work. First, due to computational constraints, our empirical evaluation is mainly limited to models at the 1.5B and 7B scales. While consistent improvements across these settings demonstrate the effectiveness of the proposed approach, experiments on larger-scale reasoning models would further strengthen the assessment of scalability. Second, our training mainly focuses on mathematical reasoning datasets, which provide reliable verifiable rewards and a broad difficulty spectrum for learning adaptive thinking. Although evaluations on GPQA and MMLU datasets indicate promising out-of-distribution generalization, extending training and evaluation to more diverse domains with trustworthy verification mechanisms remains beyond the scope of this work.

Acknowledgments

This work was supported by the Shanghai Artificial Intelligence Laboratory. We thank the anonymous reviewers for their helpful feedback.

References

- Pranjal Aggarwal and Sean Welleck. 2025. L1: Controlling how long a reasoning model thinks with reinforcement learning. *arXiv preprint arXiv:2503.04697*.
- Simon A Aytes, Jinheon Baek, and Sung Ju Hwang. 2025. Sketch-of-thought: Efficient llm reasoning with adaptive cognitive-inspired sketching. *arXiv preprint arXiv:2503.05179*.
- Qiguang Chen, Dengyun Peng, Jinhao Liu, HuiKang Su, Jiannan Guan, Libo Qin, and Wanxiang Che. 2025a. Aware first, think less: Dynamic boundary self-awareness drives extreme reasoning efficiency in large language models. *arXiv preprint arXiv:2508.11582*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025b. Do not think that much for $2+3=?$ on the overthinking of o1-like llms. *Preprint*, arXiv:2412.21187.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *Preprint*, arXiv:2110.14168.
- Muzhi Dai, Chenxu Yang, and Qingyi Si. 2025. S-grpo: Early exit via reinforcement learning in reasoning models. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pages 48178–48204. Curran Associates, Inc.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. Deepseek-v3 technical report. *Preprint*, arXiv:2412.19437.
- Ken Deng, Zizheng Zhan, Wen Xiang, Wenqiang Zhu, Weihao Li, Jingxuan Xu, Tianhao Peng, Xinping Lei, Kun Wu, Yifan Yao, Haoyang Huang, Huaixi Tang, Kepeng Lei, Zhiyi Lai, Songwei Yu, Zongxian Feng, Zuchen Gao, Weihao Xie, Chenchen Zhang, and 9 others. 2025. Hipo: Hybrid policy optimization for dynamic reasoning in llms. *Preprint*, arXiv:2509.23967.
- Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025. Thinkless: Llm learns when to think. *arXiv preprint arXiv:2505.13379*.
- Sicheng Feng, Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025. Efficient reasoning models: A survey. *arXiv preprint arXiv:2504.10903*.
- Xinshun Feng, Xinhao Song, Lijun Li, Gongshen Liu, and Jing Shao. 2026. Searl: Joint optimization of policy and tool graph memory for self-evolving agents. *arXiv preprint arXiv:2604.07791*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. *Nature*, 645(8081):633–638.
- Yang He, Xiao Ding, Bibo Cai, Yufei Zhang, Kai Xiong, Zhouhao Sun, Bing Qin, and Ting Liu. 2025a. Self-route: Automatic mode switching via capability estimation for efficient reasoning. *arXiv preprint arXiv:2505.20664*.
- Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian Yu, Zhenwen Liang, Wenxuan Wang, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025b. Deepmath-103k: A large-scale, challenging, decontaminated, and verifiable mathematical dataset for advancing reasoning. *Preprint*, arXiv:2504.11456.

- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Qitian Jason Hu, Jacob Bieker, Xiuyu Li, Nan Jiang, Benjamin Keigwin, Gaurav Ranganath, Kurt Keutzer, and Shriyash Kaustubh Upadhyay. 2024. Router-bench: A benchmark for multi-llm routing system. *arXiv preprint arXiv: 2403.12031*.
- Lingjie Jiang, Xun Wu, Shaohan Huang, Qingxiu Dong, Zewen Chi, Li Dong, Xingxing Zhang, Tengchao Lv, Lei Cui, and Furu Wei. 2025. Think only when you need with large hybrid-reasoning models. *arXiv preprint arXiv:2505.14631*.
- Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. 2024. Numinamath. [<https://huggingface.co/datasets/AI-MO/NuminaMath-1.5>](https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf).
- Guosheng Liang, Longguang Zhong, Ziyi Yang, and Xiaojun Quan. 2025. Thinkswitcher: When to think hard, when to think fast. *arXiv preprint arXiv:2505.14183*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2024. [Let's verify step by step](#). In *International Conference on Learning Representations*, volume 2024, pages 39578–39601.
- Yue Liu, Jiaying Wu, Yufei He, Ruihan Gong, Jun Xia, Liang Li, Hongcheng Gao, Hongyu Chen, Baolong Bi, Jiaheng Zhang, Zhiqi Huang, Bryan Hooi, Stan Z. Li, and Keqin Li. 2025. [Efficient inference for large reasoning models: A survey](#). *Preprint*, arXiv:2503.23077.
- Chenwei Lou, Zewei Sun, Xinnian Liang, Meng Qu, Wei Shen, Wenqi Wang, Yuntao Li, Qingping Yang, and Shuangzhi Wu. 2025. Adacot: Pareto-optimal adaptive chain-of-thought triggering via reinforcement learning. *arXiv preprint arXiv:2505.11896*.
- Anton Lozhkov, Hynek Kydlíček, Loubna Ben Allal, Guilherme Penedo, Edward Beeching, Quentin Galouédec, Nathan Habib, Lewis Tunstall, and Leandro von Werra. 2025. Openr1-math-220k. <https://huggingface.co/datasets/open-r1/OpenR1-Math-220k>.
- Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. 2025a. O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning. *arXiv preprint arXiv:2501.12570*.
- Michael Luo, Sijun Tan, Justin Wong, Xiaoxiang Shi, William Y. Tang, Manan Roongta, Colin Cai, Jeffrey Luo, Li Erran Li, Raluca Ada Popa, and Ion Stoica. 2025b. Deepscaler: Surpassing o1-preview with a 1.5b model by scaling rl. <https://pretty-radio-b75.notion.site/DeepScaleR-Surpassing-O1-Preview-with-a-1-5B-Model-by-Scaling-RL-19681902c1468005bed8ca303013a4e2>. Notion Blog.
- Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. 2025. [Cot-valve: Length-compressible chain-of-thought tuning](#). *Preprint*, arXiv:2502.09601.
- Ziqi Miao, Haonan Jia, Lijun Li, Chen Qian, Yuan Xiong, Wenting Yan, and Jing Shao. 2026. [Seeing with you: Perception-reasoning coevolution for multimodal reasoning](#). *Preprint*, arXiv:2603.28618.
- Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M Waleed Kadous, and Ion Stoica. 2025. [Routellm: Learning to route llms with preference data](#). *Preprint*, arXiv:2406.18665.
- OpenAI, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, and 243 others. 2024. [Openai o1 system card](#). *Preprint*, arXiv:2412.16720.
- Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.
- Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng, Haibin Lin, and Chuan Wu. 2025. Hybridflow: A flexible and efficient rlhf framework. In *Proceedings of the Twentieth European Conference on Computer Systems*, pages 1279–1297.
- Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, Akshay Nathan, Alan Luo, Alec Helyar, Aleksander Madry, Aleksandr Efremov, Aleksandra Spyra, Alex Baker-Whitcomb, Alex Beutel, Alex Karpenko, and 467 others. 2026. [Openai gpt-5 system card](#). *Preprint*, arXiv:2601.03267.

- Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Na Zou, Hanjie Chen, and Xia Hu. 2025. Stop overthinking: A survey on efficient reasoning for large language models. *arXiv preprint arXiv:2503.16419*.
- Yongyao Wang, Ziqi Miao, Lu Yang, Haonan Jia, Wenting Yan, Chen Qian, and Lijun Li. 2026. Tabsieve: Explicit in-table evidence selection for tabular prediction. *arXiv preprint arXiv:2602.11700*.
- Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. 2025. [TokenSkip: Controllable chain-of-thought compression in LLMs](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 3351–3363, Suzhou, China. Association for Computational Linguistics.
- Fengli Xu, Qian Yue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Chenyang Shao, Yuwei Yan, Qinglong Yang, Yiwen Song, Sijian Ren, Xinyuan Hu, Yu Li, Jie Feng, Chen Gao, and Yong Li. 2025. [Towards large reasoning models: A survey of reinforced reasoning with large language models](#). *Preprint*, arXiv:2501.09686.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. [Qwen2.5-math technical report: Toward mathematical expert model via self-improvement](#). *Preprint*, arXiv:2409.12122.
- Jingyang Yi, Jiazheng Wang, and Sida Li. 2025. [Shorterbetter: Guiding reasoning models to find optimal inference length for efficient reasoning](#). In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pages 39011–39043. Curran Associates, Inc.
- Qiyang Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaozhong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, and 16 others. 2025a. Dapo: An open-source llm reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476*.
- Tianyu Yu, Bo Ji, Shouli Wang, Shu Yao, Zefan Wang, Ganqu Cui, Lifan Yuan, Ning Ding, Yuan Yao, Zhiyuan Liu, Maosong Sun, and Tat-Seng Chua. 2025b. Rlpr: Extrapolating rlvr to general domains without verifiers. *arXiv preprint arXiv:2506.18254*.
- Linan Yue, Yichao Du, Yizhi Wang, Weibo Gao, Fangzhou Yao, Li Wang, Ye Liu, Ziyu Xu, Qi Liu, Shimin Di, and Min-Ling Zhang. 2025. [Don't overthink it: A survey of efficient rl-style large reasoning models](#). *Preprint*, arXiv:2508.02120.
- Jiajie Zhang, Nianyi Lin, Lei Hou, Ling Feng, and Juanzi Li. 2025a. [Adaptthink: Reasoning models can learn when to think](#). *arXiv preprint arXiv:2505.13417*.
- Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. 2025b. The lessons of developing process reward models in mathematical reasoning. *arXiv preprint arXiv:2501.07301*.
- Chujie Zheng, Shixuan Liu, Mingze Li, Xiong-Hui Chen, Bowen Yu, Chang Gao, Kai Dang, Yuqiong Liu, Rui Men, An Yang, Jingren Zhou, and Junyang Lin. 2025. [Group sequence policy optimization](#). *arXiv preprint arXiv:2507.18071*.

Appendix

A Experimental Details

A.1 Data Construction

To initialize both thinking and no-thinking behaviors in a unified policy, we construct a *hybrid-formatted training corpus* that explicitly contains paired supervision for the two reasoning modes.

Data Sources. The HFT dataset is built from large-scale public mathematical corpora, including OpenR1-Math, DeepMath-103K, and NuminaMath, which cover diverse problem types requiring numerical computation and multi-step reasoning.

Hybrid Supervision Construction. For each problem, we construct paired supervision corresponding to two reasoning modes:

- **Thinking mode.** Long-form trajectories are distilled from DeepSeek-R1-0528, prefixed with the control token /think and containing an explicit reasoning block <think>...</think> followed by the final answer.
- **No-thinking mode.** Short direct-answer responses are distilled from DeepSeek-V3-0324, prefixed with /no_think and including an empty reasoning block <think></think>.

The two types of supervision are balanced with a 1:1 ratio to avoid introducing bias toward either reasoning mode during supervised warm-up.

Verification. All constructed responses are verified using the Math-Verify module, and samples that fail verification are discarded. This step ensures correctness and reliability of both thinking and no-thinking supervision.

Overall, after verification and filtering, we obtain a hybrid corpus consisting of approximately 600K thinking samples and 600K no-thinking samples, providing clean, balanced supervision for initializing adaptive thinking policies.

A.2 Baselines

To comprehensively evaluate the effectiveness of the proposed approach, we compare against a diverse set of baselines that represent different strategies for controlling reasoning length and inference efficiency. All baselines are evaluated under the same experimental settings and verification protocols unless otherwise specified.

Base Variants. We first consider several variants derived from the same base language model to isolate the effect of reasoning behaviors:

- **Thinking-only.** The model always operates in the thinking mode, generating explicit chain-of-thought reasoning for every input. This variant provides the deepest level of reasoning but incurs substantial token overhead.
- **No-thinking-only.** The model always generates direct answers without intermediate reasoning. While this variant is computationally efficient, it typically suffers from degraded performance on complex problems.
- **HFT.** The hybrid fine-tuned model trained with both thinking and no-thinking supervision, as described in Appendix A.1. This model possesses both capabilities but does not autonomously decide which mode to use.

Length-Penalization Methods. We include methods that discourage excessively long reasoning chains by explicitly penalizing output length:

- **O1-Pruner** (Luo et al., 2025a). This approach introduces length-aware constraints during fine-tuning to prune redundant reasoning steps while maintaining correctness. It operates exclusively within the thinking mode and does not support explicit no-thinking behavior.
- **ShorterBetter** (Yi et al., 2025). This method encourages concise reasoning by using the shortest correct response among multiple samples as a dynamic reward signal. It guides the model toward shorter valid reasoning trajectories while preserving answer correctness.
- **S-GRPO** (Dai et al., 2025). This method reduces reasoning length by sampling a single reasoning trajectory and assigning serially decayed rewards according to intermediate exit positions. By rewarding earlier correct exits, it encourages the model to produce shorter chains of thought.

Routing-Based Methods. Routing approaches aim to reduce inference cost by dynamically selecting between different reasoning strategies:

- **RouteLLM** (Ong et al., 2025). This method trains an external router to select between a thinking model and a no-thinking model at

Table 3: Supplementary evaluation of adaptive thinking across model scales and backbone architectures. The table reports accuracy (Acc), average generation length (Len), and no-thinking ratio (Ratio_{NT}) on MATH-500, AIME-2024, and AIME-2025, covering both Qwen-2.5 series models and the Llama-3.1-8B backbone.

Method	MATH-500			AIME-2024			AIME-2025		
	Acc	Len	Ratio _{NT}	Acc	Len	Ratio _{NT}	Acc	Len	Ratio _{NT}
Qwen-2.5-1.5B									
Thinking	80.4	10441	0.0%	25.7	35829	0.0%	20.6	36977	0.0%
No-Thinking	60.4	2320	100.0%	9.1	10751	100.0%	7.5	11150	100.0%
HFT	72.3	7106	42.4%	16.8	23673	54.2%	14.5	25376	46.2%
Ours	84.6	4176	69.8%	28.1	23258	3.7%	25.2	22062	5.4%
Qwen-2.5-7B									
Thinking	90.6	8309	0.0%	65.6	29936	0.0%	59.3	35217	0.0%
No-Thinking	75.1	1890	100.0%	25.5	10464	100.0%	20.2	9921	100.0%
HFT	86.3	5619	45.2%	50.9	23484	42.5%	44.2	25017	40.4%
Ours	92.4	3403	77.6%	69.2	17762	12.5%	64.6	20060	9.6%
Llama3.1-8B									
Thinking	89.8	9046	0.0%	63.3	31294	0.0%	55.7	34801	0.0%
No-Thinking	72.1	2423	100.0%	20.6	9112	100.0%	17.5	10905	100.0%
HFT	82.2	5619	48.8%	44.7	22388	46.3%	37.1	22938	49.2%
Ours	90.8	3618	75.5%	66.7	19089	10.1%	60.4	21084	8.7%

inference time. While effective in reducing token usage, it relies on additional supervision for router training and requires maintaining multiple models during deployment.

Adaptive Thinking Methods. We further compare against recent reinforcement-learning-based adaptive thinking approaches:

- **ThinkLess** (Fang et al., 2025). This method adopts a decoupled GRPO formulation to encourage shorter chains, learning a hybrid policy over thinking and no-thinking behaviors.
- **AdaptThink** (Zhang et al., 2025a). This approach employs a PPO-based objective that explicitly promotes concise reasoning by increasing the likelihood of short chains.

Both ThinkLess and AdaptThink share the same HFT initialization as our method, ensuring fair comparison and controlled base model capacity.

Evaluation Protocol. All methods are evaluated using the same datasets, verification module (Math-Verify³), and decoding settings. We report accuracy, average generation length, and the no-thinking ratio to jointly measure reasoning performance and efficiency. For MATH-500, all reported results are averaged over 4 independent runs. For

AIME-2024 and AIME-2025, results are averaged over 32 independent runs to reduce variance.

A.3 Implementation Details

All experiments are conducted on 16 NVIDIA H100 GPUs. We implement the Hybrid Fine-Tuning (HFT) stage using LLaMA-Factory⁴ and the reinforcement learning stage using VeRL⁵.

HFT. In the HFT stage, the maximum context length is set to 16K tokens, with overlong samples truncated. Models are trained for 3 epochs with the AdamW optimizer, using a 10% linear warmup followed by cosine learning rate decay, and a peak learning rate of 1×10^{-4} .

Reinforcement Learning. In the RL stage, the context length is extended to 24K tokens. The policy model is optimized with AdamW using a constant learning rate of 1×10^{-6} . We use a batch size of 128 with a micro-batch size of 64, and sample $G=8$ responses per query during rollout.

Configuration Fairness. For fair comparison, all baselines under the same model scale and backbone are initialized from the same base checkpoint and trained with unified data, context-length, and optimization budgets whenever applicable. We use

³<https://github.com/huggingface/Math-Verify>

⁴<https://github.com/hiyouga/LlamaFactory>

⁵<https://github.com/volcengine/verl>

Table 4: Detailed supplementary results on extended benchmarks and complete baseline comparisons. The table reports additional results on GSM8K and AMC, together with the full baseline suite for Qwen-2.5-7B and Llama-3.1-8B. *Len* denotes the average number of generated tokens per example. Best and second-best results among non-reference methods are marked in bold and underlined, respectively.

Method	GSM8K		AMC		MATH-500		AIME-2024		AIME-2025	
	Acc	Len	Acc	Len	Acc	Len	Acc	Len	Acc	Len
Qwen-2.5-7B										
Thinking	92.3	3974	89.8	15476	90.6	8309	65.6	29936	59.3	35217
No-Thinking	84.9	683	56.5	5303	75.1	1890	25.5	10464	20.2	9921
HFT	88.7	2271	72.7	11062	86.3	5619	50.9	23484	44.2	25017
O1-Pruner	89.2	2597	87.3	12936	87.9	5567	62.4	25977	57.5	32451
ShorterBetter	89.5	2033	86.9	10374	88.1	3594	61.9	18615	57.1	21733
S-GRPO	91.1	2165	89.4	11652	90.9	5326	65.1	22893	59.7	28474
RouteLLM	90.4	2268	87.5	9974	88.2	3121	55.3	24011	51.7	28387
Thinkless	91.9	2362	89.3	9658	89.8	<u>3302</u>	64.0	18172	<u>60.1</u>	<u>21463</u>
AdaptThink	<u>92.6</u>	2419	<u>90.2</u>	<u>9563</u>	<u>91.1</u>	3462	<u>65.2</u>	<u>18123</u>	59.1	22544
Ours	94.2	<u>2064</u>	91.1	9371	92.4	3403	69.2	17762	64.6	20060
Llama-3.1-8B										
Thinking	90.9	4173	86.7	16251	89.8	9046	63.3	31294	55.7	34801
No-Thinking	83.6	717	54.2	5568	72.1	2423	20.6	9112	17.5	10905
HFT	87.4	2385	69.8	11615	82.2	5619	44.7	22388	37.1	22938
O1-Pruner	87.9	2727	83.1	13584	85.4	5171	60.1	25133	53.1	28894
ShorterBetter	88.2	2135	83.4	10893	85.6	3984	59.5	19008	54.3	22033
S-GRPO	89.7	2273	85.8	12235	88.3	5404	62.4	22147	56.4	28367
RouteLLM	89.0	2381	84.0	10473	85.7	3460	58.9	21228	49.3	24719
Thinkless	90.5	2480	85.7	10162	88.5	3838	63.5	19532	56.9	22855
AdaptThink	<u>91.2</u>	2344	<u>86.6</u>	<u>10041</u>	<u>89.4</u>	3660	<u>64.4</u>	19580	<u>57.8</u>	<u>21759</u>
Ours	92.8	<u>2261</u>	89.5	9840	90.8	<u>3618</u>	66.7	<u>19089</u>	60.4	21084

a 16K context length for HFT and a 24K context length for all RL-based methods to reduce excessive rollout truncation and avoid sparse or unreliable reward signals. During evaluation, all methods share the same prompts, decoding protocol, answer extraction procedure, and metric computation.

B Additional Experimental Results

To further assess the generality of our approach, we conduct additional experiments beyond the Qwen2.5 series. Specifically, we apply the same training pipeline to LLaMA 3.1-8B Base, which adopts a different architecture and pretraining recipe. All training and optimization procedures strictly follow the HFT and RL settings described in Appendix A.3, with only minimal hyperparameter adjustments for model size. No architecture-specific modifications are introduced.

As shown in Table 3, our method consistently improves reasoning accuracy while substantially reducing the average generation length under the LLaMA architecture. These results demonstrate that the proposed correctness-preserving advantage

shaping and length-aware gradient regulation are not specific to a particular backbone, but instead generalize well across different model families.

To further supplement the main evaluation, Table 4 reports extended results on GSM8K and AMC, together with the complete baseline suite for Qwen-2.5-7B and LLaMA 3.1-8B. Across all five benchmarks and both backbones, our method consistently achieves the best accuracy while substantially reducing generation length compared with the always-Thinking baseline. The results on GSM8K and AMC further show that the learned adaptive thinking policy is not limited to competition-level benchmarks, but can transfer to tasks with different reasoning formats and difficulty distributions. The full LLaMA baseline comparison also confirms that the observed gains remain stable when compared against competitive baselines under a different backbone. These results further support the cross-dataset robustness and cross-backbone generality of our approach.

C Further Analysis of CPAS and LAGR

This section provides a concise analysis of CPAS and LAGR in adaptive thinking. Our goal is not to establish a general convergence guarantee, but to clarify how the two mechanisms separate correctness from efficiency and regulate optimization under heterogeneous reasoning trajectories.

Advantage Suppression with Length-Dependent Rewards. In GRPO, multiple responses are sampled for the same query, and their scalar rewards are normalized within the rollout group. Several existing efficient-reasoning methods directly incorporate an efficiency preference into the scalar reward. To distinguish this setting from the reward used in our method, let

$$\tilde{r}_i = R_{\text{acc}}(o_i) + \gamma R_{\text{len}}(o_i)$$

denote a reward that directly combines answer correctness and response length. The corresponding group-relative advantage is

$$\tilde{A}_i^{\text{GRPO}} = \frac{\tilde{r}_i - \text{Mean}(\tilde{r})}{\text{Std}(\tilde{r})}.$$

When shorter correct trajectories systematically receive higher values of $\tilde{r}_i$, a longer but still correct trajectory may satisfy

$$0 < \tilde{r}_i < \text{Mean}(\tilde{r}),$$

and consequently receive $\tilde{A}_i^{\text{GRPO}} < 0$. In this case, the negative advantage is caused not by answer incorrectness, but by the relative comparison induced by the length-dependent reward. Such a signal can suppress valid long-form reasoning and bias the policy toward shorter responses.

Correctness-Only Base Reward. Our method avoids directly coupling correctness and length in the scalar rollout reward. Specifically, the base reward used to compute the GRPO advantage is determined solely by answer correctness:

$$r_i = R_{\text{acc}}(o_i) \in \{0, 1\}.$$

The corresponding base advantage is

$$A_i^{\text{GRPO}} = \frac{r_i - \text{Mean}(r)}{\text{Std}(r)}.$$

Therefore, correct direct-answer responses and correct long-form reasoning responses receive the same base reward before advantage shaping.

For a rollout group containing both correct and incorrect responses, suppose that k out of G responses are correct, where $0 < k < G$. For every correct response,

$$r_i - \text{Mean}(r) = 1 - \frac{k}{G} > 0.$$

Thus, its base group-relative advantage is positive and independent of its response length. When all responses in a group have the same correctness reward, the centered base advantage provides no relative correctness signal; importantly, the responses are still not ranked according to their lengths. Consequently, a correct long-form response is not assigned a lower base reward or base advantage merely because it is longer.

Correctness-Preserving Advantage Shaping.

After computing the GRPO advantage using the correctness-only reward, CPAS introduces an additional efficiency incentive for correct short-chain responses:

$$A_i = \begin{cases} A_i^{\text{GRPO}} + \delta, & r_i = 1 \text{ and } o_i \text{ is short-chain,} \\ A_i^{\text{GRPO}}, & \text{otherwise.} \end{cases}$$

Here, $\delta > 0$ is a small coefficient that encourages efficient responses only when the final answer is correct.

In this paper, “short-chain” refers specifically to a response generated under the direct-answer mode. It is determined by the mode selected at the initial mode-decision position, rather than by a manually specified token-length threshold or by a relative comparison with other responses in the rollout group. We use the term “short-chain” because this mode generally produces substantially shorter responses.

This design separates correctness from efficiency. The base GRPO advantage is computed solely from answer correctness, so correct short-chain and long-chain responses receive the same base reward before advantage shaping. CPAS then provides an additional advantage only to correct short-chain responses. Incorrect short-chain responses do not receive this bonus, while correct long-chain responses retain their original correctness-based GRPO advantage without an explicit length-dependent penalty. Therefore, CPAS acts as a correctness-gated efficiency incentive rather than compensating for a long-chain response after it has received a negative length-dependent advantage.

Table 5: Comparison with alternative length-normalization schemes on Qwen2.5-7B. Acc denotes accuracy, Len denotes the average number of generated tokens, and Ratio_{NT} denotes the proportion of responses generated under the direct-answer mode. All reinforcement-learning methods use the same initialization and optimization settings.

Method	MATH-500			AIME-2024			AIME-2025		
	Acc	Len	Ratio _{NT}	Acc	Len	Ratio _{NT}	Acc	Len	Ratio _{NT}
HFT	86.3	5619	45.2%	50.9	23484	42.5%	44.2	25017	40.4%
DAPO	91.7	4761	63.4%	69.1	22168	6.7%	64.9	22495	3.5%
Dr. GRPO	92.1	4957	59.7%	68.8	24169	6.5%	64.3	23841	3.1%
LAGR (Ours)	92.4	3403	77.6%	69.2	17762	12.5%	64.6	20060	9.6%

Length-Aware Gradient Regulation. Besides the interaction between correctness and efficiency at the advantage level, adaptive thinking introduces gradient imbalance due to response-length heterogeneity. Under response-level GRPO averaging, the gradient contribution of response o_i can be written as

$$g_i^{\text{GRPO}} = \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \hat{A}_{i,t} \nabla_{\theta} \log \pi_{\theta}(o_{i,t} \mid q, o_{i,<t}).$$

The factor $1/|o_i|$ averages token-level gradients within each response. When a rollout group contains responses with highly heterogeneous lengths, this normalization can reduce the effective contribution of long responses and make optimization relatively more sensitive to short ones.

LAGR regulates this imbalance by reweighting response-level gradient contributions according to response length:

$$g_i^{\text{LAGR}} = w_i^{(\beta)} \sum_{t=1}^{|o_i|} \hat{A}_{i,t} \nabla_{\theta} \log \pi_{\theta}(o_{i,t} \mid q, o_{i,<t}),$$

where

$$w_i^{(\beta)} = \frac{1}{M} \cdot \frac{|o_i|^{-\beta}}{\sum_{j=1}^G |o_j|^{-\beta}}, \quad \beta \in [0, 1].$$

The coefficient β controls the strength of length-aware regulation. When $\beta = 1$, the update remains close to the original length-sensitive response-averaging behavior. When $\beta = 0$, responses receive uniform response-level weights, increasing the relative contribution of long trajectories. Intermediate values continuously regulate gradient allocation between these two endpoints, providing a controllable balance between efficient direct answering and preservation of long-form reasoning.

Comparison with Alternative Length-Normalization Schemes. We compare LAGR with the length-normalization schemes used in DAPO and Dr. GRPO. DAPO adopts token-level loss aggregation, while Dr. GRPO uses length-independent normalization. For a controlled comparison, all methods share the same HFT initialization, training data, rollout budget, optimization settings, correctness reward, and evaluation protocol.

As shown in Table 5, both alternatives improve substantially over the HFT initialization, confirming the importance of gradient allocation under heterogeneous response lengths. Compared with them, LAGR consistently generates shorter responses while achieving comparable or better accuracy across the three benchmarks. These results suggest that LAGR’s continuous regulation of length-dependent gradient contributions provides a more favorable accuracy–efficiency trade-off than fixed aggregation schemes.

Connection to Empirical Analysis. The above analysis clarifies the distinct roles of CPAS and LAGR. CPAS separates correctness and efficiency by computing the base advantage solely from answer correctness and applying an additional efficiency incentive only to correct direct-answer responses. LAGR operates at the gradient-allocation level and regulates the contribution of responses with substantially different lengths. Their practical effects are further examined through the ablation and sensitivity studies in Section 4.4, including CPAS-related training dynamics and the influence of the LAGR parameters β and λ . Together, these mechanisms support a more stable accuracy–efficiency trade-off under heterogeneous reasoning trajectories.

D Case Study

Question (MATH500): Convert the point $(0, 3)$ in rectangular coordinates to polar coordinates. Enter your answer in the form (r, θ) , where $r > 0$ and $0 \leq \theta < 2\pi$. Format your response as follows: “The correct answer is boxed{insert answer here}.”

Thinking:

<think>I need to convert the point $(0, 3)$ from rectangular coordinates to polar coordinates. Polar coordinates are given as (r, θ) , where r is the distance from the origin to the point, and θ is the angle from the positive x -axis. First, recall the formulas for converting from rectangular to polar coordinates. The formulas are: $r = \sqrt{x^2 + y^2}$ and $\theta = \tan^{-1}(y/x)$. [...] **(omit)** which is not the same point. So, only $\theta = \pi/2$ is correct. I could use $\theta = \pi/2 + 2k\pi$, but within 0 to 2π , only $\pi/2$. 2π would be $5\pi/2$, which is greater than 2π , so no. So, I think $(3, \pi/2)$ is the answer. The problem says “enter your answer in the form (r, θ) ”, and “boxed”. So, I need to write it as $(3, \frac{\pi}{2})$ in the box. The instruction says: “output the final answer within \boxed{ }” and “enter your answer as (r, θ) ”, so, probably $(3, \frac{\pi}{2})$. But let me make sure the format. Sometimes they write it as a pair. I think that’s fine. Is there any other representation? For example, could r be negative? But the problem specifies $r > 0$, so no. Or θ could be different, but no. Another thing: when $x = 0$ and $y > 0$, it’s unambiguously $\pi/2$. If it were $(0, -3)$, it would be $3\pi/2$ or $-\pi/2$, but since it’s positive, it’s fine. So, I think that’s it. The point is $(0, 3)$, which is on the y -axis, so $\theta = 90$ degrees, which is $\pi/2$ radians. Yes. So, final answer. **</think>** Final Answer: To convert the point $(0, 3)$ from rectangular coordinates to polar coordinates (r, θ) with $r > 0$ and $0 \leq \theta < 2\pi$, calculate $r = \sqrt{x^2 + y^2} = \sqrt{0^2 + 3^2} = 3$. Since $x = 0$ and $y > 0$, the point lies on the positive y -axis, so $\theta = \pi/2$, and $\pi/2 \in [0, 2\pi)$. Thus, the polar coordinates are $(3, \pi/2)$, i.e., the correct answer is $\boxed{(3, \frac{\pi}{2})}$. **(3272 tokens)**

Ours:

/nothink <think></think> To convert the rectangular coordinates $((0, 3))$ to polar coordinates (r, θ) , follow these steps: 1. Calculate r (the distance from the origin to the point): The formula for r is: $r = \sqrt{x^2 + y^2}$ Substituting $x = 0$ and $y = 3$: $r = \sqrt{0^2 + 3^2} = \sqrt{9} = 3$ 2. Determine θ (the angle with the positive x -axis): Since the point $((0, 3))$ lies on the positive y -axis, the angle θ is: $\theta = \frac{\pi}{2}$ This is because the positive y -axis corresponds to 90° or $(\frac{\pi}{2})$ radians in the range $0 \leq \theta < 2\pi$. 3. Combine the results: The polar coordinates are: $(3, \frac{\pi}{2})$ Final Answer: $\boxed{(3, \frac{\pi}{2})}$ **(713 tokens)**

Figure 5: For an example from MATH-500, the Thinking baseline generates a long chain-of-thought with redundant intermediate steps. In contrast, our method chooses *NoThinking* and directly produces a concise final solution, using 713 tokens in total.

Question (AIME): Find the largest possible real part of $(75 + 117i)z + \frac{96 + 144i}{z}$ where z is a complex number with $|z| = 4$.

⚙️ Thinking:

<think> I need to find the largest possible real part of $(75 + 117i)z + \frac{96+144i}{z}$ where z is a complex number with magnitude 4. Since $|z| = 4$, I can write $z = 4e^{i\theta}$ for some real θ , or equivalently, $z = 4(\cos \theta + i \sin \theta)$. Let me denote the expression as $w = (75 + 117i)z + \frac{96+144i}{z}$. I need the real part of w , and maximize it. Since z is on the circle of radius 4, I could set $z = x + iy$ with $x^2 + y^2 = 16$, but that might get messy with the algebra. Notice that the expression involves both z and $1/z$, and since $|z| = 4$, $|1/z| = 1/4$. Also, the coefficients are complex numbers, so I need to handle that. Let me simplify the coefficients. [...] **(omit)** So real part is $\text{Re}((r + is)/4 * (\cos \theta - i \sin \theta)) = (1/4) \text{Re}((r + is)(\cos \theta - i \sin \theta)) = (1/4) [r \cos \theta - r i \sin \theta + i s \cos \theta - i^2 s \sin \theta] = (1/4) [r \cos \theta + s \sin \theta + i(-r \sin \theta + s \cos \theta)]$ So real part is $(1/4)(r \cos \theta + s \sin \theta)$ Similarly, from first term: $4[(p + iq)(\cos \theta + i \sin \theta)] = 4[p \cos \theta - q \sin \theta + i(p \sin \theta + q \cos \theta)]$, so real part $4(p \cos \theta - q \sin \theta)$ Thus total real part $\text{Re}(w) = 4(p \cos \theta - q \sin \theta) + \frac{1}{4}(r \cos \theta + s \sin \theta)$ Now plug in $p=75, q=117, r=96, s=144$. So $\text{Re}(w) = 4(75 \cos \theta - 117 \sin \theta) + \frac{1}{4}(96 \cos \theta + 144 \sin \theta)$ Compute: $4*75=300, 4*(-117)=-468$, so $300 \cos \theta - 468 \sin \theta$ Then $\frac{1}{4} * 96 = 24, \frac{1}{4} * 144 = 36$, so $+24 \cos \theta + 36 \sin \theta$ Thus total: $(300 + 24) \cos \theta + (-468 + 36) \sin \theta = 324 \cos \theta - 432 \sin \theta$, same as before. So no mistake. Thus maximum is $\sqrt{324^2 + (-432)^2} = \sqrt{104976 + 186624} = \sqrt{291600} = 540$. So the largest possible real part is 540. Therefore, the answer is 540. But the problem says "real part", and I got it, so I think that's it. **</think>** To find the largest possible real part of $(75 + 117i)z + \frac{96+144i}{z}$ where z is a complex number with $|z| = 4$, express z in polar form as $z = 4e^{i\theta}$ for some real θ . This gives $z = 4(\cos \theta + i \sin \theta)$. Define the expression as $w = (75 + 117i)z + \frac{96+144i}{z}$. Substitute z : $w = (75 + 117i) \cdot 4(\cos \theta + i \sin \theta) + \frac{96+144i}{4(\cos \theta + i \sin \theta)}$. Simplify the second term: $\frac{96+144i}{4(\cos \theta + i \sin \theta)} = \frac{96+144i}{4} \cdot \frac{1}{\cos \theta + i \sin \theta} = (24 + 36i) \cdot \frac{\cos \theta - i \sin \theta}{1}$, since $\frac{1}{\cos \theta + i \sin \theta} = \cos \theta - i \sin \theta$ for $|z| = 4$. [...] **(omit)** This maximum is achieved for some θ , such as when $\cos \theta = \frac{3}{5}$ and $\sin \theta = -\frac{4}{5}$, giving $\text{Re}(w) = 540$. Since the function is continuous and the circle $|z| = 4$ is compact, the maximum value is attained. Thus, the largest possible real part is 540. **(22383 tokens)**

⚙️ Ours:

/nothink <think></think> Alright, I have this problem: Find the largest possible real part of $(75 + 117i)z + \frac{96+144i}{z}$ where z is a complex number with $|z| = 4$. Understanding the Problem First, I need to understand what's being asked. We have a complex expression in terms of z , and z is a complex number with magnitude 4. The goal is to find the maximum real part of this expression. Step 1: Represent z in a Convenient Form Since z is a complex number with $|z| = 4$, it can be written in polar form as: $z = 4e^{i\theta}$ where θ is the argument of z . But before jumping into polar form, maybe it's easier to think of z in terms of its real and imaginary parts: $z = x + iy$ where x, y are real numbers, and $x^2 + y^2 = 16$ because $|z| = 4$. Step 2: Rewrite the Expression The expression is: [...] **(omit)** $x = \frac{3}{5} \times 4 = \frac{12}{5}, y = \frac{-4}{5} \times 4 = \frac{-16}{5}$ Now, plug these into $(81x - 108y)$: $81 \times \frac{12}{5} - 108 \times \frac{-16}{5} = \frac{972}{5} + \frac{1728}{5} = \frac{972+1728}{5} = \frac{2700}{5} = 540$ This matches our previous result. Conclusion The largest possible real part of the given expression is 540. **(2977 tokens)**

Figure 6: For an example from AIME, the Thinking baseline generates a long chain-of-thought with redundant intermediate steps. In contrast, our method chooses *NoThinking* and directly produces a concise final solution, using 2977 tokens in total.

🧑 Question (GPQA): Answer the following multiple choice question. The last line of your response should be of the following format: 'Answer: \$LETTER' (without quotes) where LETTER is one of ABCD. Think step by step before answering. "Oh, I know you," the ribonucleoprotein particle says to the nascent chain as they meet. "Pause there for a minute. Let me show you in; you really need some sugar." "It seems somewhat rough. I guess this is goodbye; I need to be on my way", the chain replies. Where did they meet, and where is the chain heading? A) Golgi to the mitochondrion B) membrane to the nucleus C) cytosol to the extracellular space D) ribosome to the proteasome.

⚙️ Thinking:

<think> I need to figure out where the ribonucleoprotein particle and the nascent chain meet, and where the chain is heading. The options are: A) Golgi to the mitochondrion B) membrane to the nucleus C) ribosome to the proteasome D) cytosol to the extracellular space First, I recall that the ribonucleoprotein particle (RNP) is involved in transporting nucleotides into cells. Nucleotides are the building blocks of DNA and RNA. The nascent chain is probably a nascent protein, which is a protein that's being synthesized by the ribosome. [...] **(omit)** So, option D. Perhaps in some contexts, "cytosol" is used for the cell membrane. I think I'll go with D. But let's see if there's a better fit. Option A: "Golgi to the mitochondrion" – if the chain is heading to the Golgi, but why would it be at "Golgi to mitochondrion". Perhaps for a protein that is going to the mitochondrion, but the question is about a general nascent chain needing nucleotides, not necessarily for a specific location. So, unlikely. Perhaps in this case, after meeting, it's heading to the Golgi, and the meeting is at the cytoplasm, but not specified. I think D is the most plausible. Perhaps B: "membrane to the nucleus" – for a protein that is being integrated, but again, not for nucleotide entry. I think D is the answer. So, I'll say D. But let's confirm. Final answer. **</think>** The ribonucleoprotein particle (RNP) shuttles nucleotides (Sugar) into cells, and the nascent chain (a protein being synthesized) is present in the cytoplasm. [...] **(omit)** - B) Membrane to the nucleus: Incorrect, as this is typically for secreted proteins that integrate into the nucleus, but the chain needs nucleotides for assembly, not nuclear integration. - C) Ribosome to the proteasome: Incorrect, as the ribosome and proteasome are locations for different processes; the ribosome assembles the protein, while proteasomal degradation occurs later. - D) Cytosol to extracellular space: Correct, as the meeting occurs at the cell membrane (boundary of cytoplasm and extracellular space), and the chain heads to the extracellular space after receiving the nucleotides. Thus, they meet at the cell membrane (described as "cytosol to extracellular space"), and the chain is heading to the extracellular space. **(22149 tokens)**

⚙️ Ours:

/nothink <think></think> Let's think step by step to answer the question. Understanding the Question: The question is a multiple-choice question about the meeting place and destination of a nascent (newly synthesized) protein chain. The options are: - A) Golgi to the mitochondrion [...] **(omit)** This fits because the cytoplasm (where the ribonucleoprotein is located) processes the protein, and the processed protein then exits the cell to the extracellular space. - D) Ribosome to the proteasome: The ribosome is where protein synthesis occurs, and the proteasome is involved in protein degradation. The described process is not about moving from the ribosome to the proteasome but about processing and exiting. Conclusion: The correct answer is C) cytosol to the extracellular space, as the ribonucleoprotein processes the nascent protein in the cytoplasm, and the processed protein is then released into the extracellular space. **(2197 tokens)**

Figure 7: For an example from GPQA, the Thinking baseline generates a long chain-of-thought with redundant intermediate steps. In contrast, our method chooses *NoThinking* and directly produces a concise final solution, using 2197 tokens in total.