

# Reasoning Pattern Alignment Merging for Adaptive Reasoning

Anonymous ACL submission

## Abstract

Recent large reasoning models (LRMs) have made substantial progress in complex reasoning tasks, yet they often generate lengthy reasoning paths for every query, incurring unnecessary computation and latency. Existing speed-up approaches typically rely on retraining the model or designing sophisticated prompting, which are either prohibitively expensive or highly sensitive to the input and prompt formulation. In this work, we study model merging as a lightweight alternative for efficient reasoning: by combining a long chain-of-thought (Long-CoT) reasoning model with a Short-CoT instruction model, we obtain an adaptive reasoner without training from scratch or requiring large-scale additional data. Building on this idea, we propose Reasoning Pattern Alignment Merging (RPAM), a layer-wise model merging framework based on feature alignment to facilitate query-adaptive reasoning. RPAM first constructs a small pattern-labeled calibration set that assigns each query an appropriate reasoning pattern. It then optimizes layer-wise merging coefficients by aligning the merged model's intermediate representations with those of the selected model, while a contrastive objective explicitly pushes them away from the non-selected model. Experiments on seven widely used reasoning benchmarks show that RPAM substantially reduces inference cost while maintaining strong performance.

## 1 Introduction

Large reasoning models (LRMs) have recently achieved strong performance on complex reasoning tasks, ranging from mathematical problem solving (Lewkowycz et al., 2022; Pan et al., 2024; Shao et al., 2024) and logical deduction (Wei et al., 2022) to agentic assistants (Wu et al., 2025b; Chen et al., 2025a). A key driver is their ability to generate long chains of thought (Long-CoT), in which the model iteratively self-assesses, mitigates errors, and verifies intermediate steps before producing a final

answer (Liu et al., 2024; Yu et al., 2024b; Li et al., 2025). However, while Long-CoT is beneficial for difficult problems, it can be counterproductive on simple tasks that require few reasoning steps: models may "overthink" by introducing unnecessary intermediate reasoning (Muennighoff et al., 2025; Zhang et al., 2025b). This overthinking not only increases inference cost (Fan et al., 2025; Chen et al., 2025b; Cuadron et al., 2025), but can also hurt accuracy by amplifying the chance of spurious reasoning and obscuring the straightforward solution (Shen et al., 2025).

Many recent studies have explored how to mitigate the inefficiency and overthinking of LRMs (Feng et al., 2025; Srivastava et al., 2025; Sui et al., 2025). A prominent direction is adaptive reasoning, where a model dynamically selects between Long-CoT and Short-CoT modes depending on the estimated problem complexity (Zhang et al., 2025a; Alomrani et al., 2025). Broadly, existing adaptive-reasoning approaches fall into two paradigms. Training-based methods optimize LRMs to exhibit adaptive thought processes via supervised fine-tuning (SFT) (Yu et al., 2025; Ma et al., 2025) or reinforcement learning (RL) (Fang et al., 2025; Wu et al., 2025c), but they typically require large-scale data and incur substantial training cost. Training-free methods, often implemented through prompt-guided strategies (Xu et al., 2025b), can introduce adaptivity without additional optimization, yet they rely heavily on instruction following and can be sensitive to prompt constraints and query phrasing (Zhu et al., 2025; Zhu and Li, 2025). As a result, achieving adaptive reasoning that is both efficient and effective remains non-trivial.

In light of this, we investigate model merging as an alternative route toward adaptive reasoning. Model merging can integrate complementary capabilities from multiple task- or style-specific models into a single model without training from scratch

and often without requiring large-scale additional data (Wortsman et al., 2022; Ilharco et al., 2023). This property makes it a natural fit for adaptive reasoning, since Long-CoT and Short-CoT behaviors are commonly learned in separate models. Nevertheless, most prior merging work primarily targets improved overall accuracy of a single merged model (Wu et al., 2025a), rather than explicitly optimizing for query-adaptive reasoning efficiency. Recent work, such as ACM (Yao et al., 2025), moves toward efficiency by shortening reasoning trajectories globally across queries, which may be suboptimal for inputs that genuinely require long-horizon deliberation.

Motivated by this gap, we aim for *query-level* adaptive reasoning via model merging: the merged model should produce long, detailed reasoning when necessary, but remain concise on simple inputs. To this end, we propose *Reasoning Pattern Alignment Merging* (RPAM), a progressive, layer-wise merging framework with a contrastive objective. Starting from a Long-CoT reasoning model and a Short-CoT instruction model, we first construct a small pattern-labeled dataset in which each query is automatically assigned a reasoning-pattern label (Long-CoT vs. Short-CoT) based on which model achieves higher expected correctness. We then optimize layer-wise merging coefficients so that the merged model’s intermediate features align with those of the selected (positive) model for each instance, while a contrastive loss explicitly pushes them away from the non-selected (negative) model. This yields a merged model that preserves strong performance on challenging problems while avoiding unnecessary long reasoning on simple ones, reducing inference cost without sacrificing accuracy. We evaluate our approach across seven commonly used reasoning benchmarks. Empirical results show that our method reduces token generation by 48% on the MATH dataset (Hendrycks et al., 2021) while simultaneously improving accuracy. Furthermore, on the more challenging OlympiadBench (He et al., 2024), our approach achieves a 50% reduction in inference cost with only negligible accuracy degradation.

Our contributions are summarized as follows:

- We introduce a model-merging-based paradigm for query-level adaptive reasoning that balances reasoning depth and efficiency within a single merged model.
- We develop a simple yet effective merging

recipe, a small pattern-labeled dataset set together with feature alignment and contrastive separation, to learn layer-wise merging coefficients that integrate Long-CoT and Short-CoT behaviors.

- Through extensive experiments, we demonstrate that RPAM achieves a stronger accuracy and efficiency trade-off than competitive model-merging and adaptive-reasoning baselines, attaining near Long-CoT performance while substantially reducing generated tokens.

## 2 Related Work

### 2.1 Adaptive Reasoning

Adaptive reasoning improves efficiency by dynamically selecting an appropriate reasoning pattern conditioned on query difficulty (Sui et al., 2025; Zhu and Li, 2025). Existing approaches broadly fall into two lines: prompt-guided and training-based methods. **Prompt-guided** approaches (Renze and Guven, 2024; Gong et al., 2025; Xu et al., 2025b) exploit the instruction-following capability of LRMs by imposing explicit constraints through carefully crafted prompts to elicit concise reasoning paths. For instance, CoUT (Gong et al., 2025) first guides LRMs to internalize thinking processes and then generates the final answer using several token-efficient decoding strategies. However, their effectiveness can be brittle, as it depends heavily on prompt design, constraint choices, and the model’s instruction-following behavior. **Training-based** methods explicitly optimize models to control or prune response length, e.g., by fine-tuning LRMs on variable-length CoT supervision (Luo et al., 2025; Qiao et al., 2025) or by applying reinforcement learning with length-aware rewards (Zhang et al., 2025a; Arora and Zanette, 2025). For example, Arora and Zanette (2025) post-train LRMs with RL using a length-penalty reward that combines correctness and normalized length of sampled solutions, assigning higher reward scores to shorter correct solutions. While effective, these approaches typically require substantial training data and incur high computational cost.

### 2.2 Model Merging

Model merging (Yang et al., 2024a) is an emerging technique that combines parameters from multiple pretrained or fine-tuned models with complementary capabilities into a single model, typically with-

out access to the original training data and without expensive end-to-end retraining. Early work such as Model Soups (Wortsman et al., 2022) demonstrates that simple weight averaging over multiple checkpoints can improve overall performance. Beyond averaging, more advanced methods, including TIES-Merging (Yadav et al., 2023) and DARE (Yu et al., 2024a), aim to reduce interference among task vectors by selectively retaining and combining only the most salient parameter updates. Building on these observations, representation surgery (Yang et al., 2024b) shows that naive merging can induce representation bias and proposes post-merging calibration by aligning the merged model’s representations with those of task-specific models.

More recently, several studies (Wu et al., 2025a; Yao et al., 2025; Wu et al., 2025d) have explored model merging for efficient reasoning by combining a slow-thinking Long-CoT model with a fast-thinking Short-CoT model. These approaches typically aim to produce a single merged model that globally shortens reasoning paths while retaining accuracy. In contrast, our work focuses on query-level adaptive reasoning: we calibrate the merged model to selectively follow Long-CoT or Short-CoT patterns depending on the input, thereby improving the accuracy and efficiency trade-off.

### 3 Preliminary

#### 3.1 Problem Formulation

Given a problem query  $x$ , an LLM parameterized by  $\theta$  generates a corresponding solution  $y = [y^1, \dots, y^k]$  by sampling from the conditional distribution  $\pi_\theta(\cdot|x)$ . This distribution factorizes autoregressively as:

$$\pi_\theta(y|x) = \prod_{i=1}^k \pi_\theta(y^i | x, y^{<i}). \quad (1)$$

To improve the solution quality on reasoning tasks, Chain-of-Thought (CoT) prompting has been widely adopted, as it encourages models to produce intermediate steps that facilitate self-evaluation and verification. According to the granularity of the generated steps, CoT reasoning can be classified into Long-CoT and Short-CoT (Luo et al., 2025). Long-CoT yields detailed, reflective reasoning that benefits complex queries but increases generation length and latency. Short-CoT, in contrast, produces concise reasoning (or even direct answers) with lower cost, but can struggle on harder problems that require multi-step deliberation.

In this paper, we consider two LLMs: a Long-CoT reasoning model  $\theta_L$  and a Short-CoT instruction model  $\theta_S$ . Our goal is to construct a new model  $\theta_M$  that can adapt its reasoning pattern to the query, producing Long-CoT when deeper reasoning is necessary and Short-CoT when a concise response suffices.

#### 3.2 Model Merging

Model merging aims to combine  $N$  well-trained neural network models  $\{\theta_i\}_{i=1}^N$ , typically specialized for different tasks or behaviors, into a single parameter set  $\theta_M$  without additional end-to-end retraining, while preserving the strengths of the constituent models. A simple and widely used approach is linear weight merging:

$$\theta_M = \sum_{i=1}^N \lambda_i \theta_i, \quad (2)$$

where  $\lambda_i$  controls the contribution of model  $\theta_i$  (often with  $\sum_i \lambda_i = 1$ ). However, because the coefficients are usually pre-defined, naive linear merging can introduce representation bias and yield suboptimal performance (Yang et al., 2024b). To alleviate this issue, post-calibration merging methods have been proposed (Dai et al., 2025; Xu et al., 2025a), which make the merging coefficients learnable and calibrate the merged model toward task-specific optima, thereby reducing knowledge loss.

In this work, we view the Long-CoT model and the Short-CoT model as two complementary, task-specialized models, and design a post-calibration model merging method to construct a single merged model  $\theta_M$  that can adapt its reasoning behavior to the input query.

### 4 Proposed Method

In this paper, we propose RPAM, a reasoning-pattern-aware model merging framework that integrates the complementary behaviors of a Long-CoT reasoning model and a Short-CoT instruction model. The goal is to obtain a single merged model that can adapt its reasoning style to the input: using deeper, longer reasoning when necessary, while remaining concise and efficient on simpler queries. As illustrated in Fig. 1, RPAM consists of two components: (i) pattern-labeled data construction, (ii) feature-alignment merging with contrastive enhanced shaping. In the following parts, we present technical details for each design.

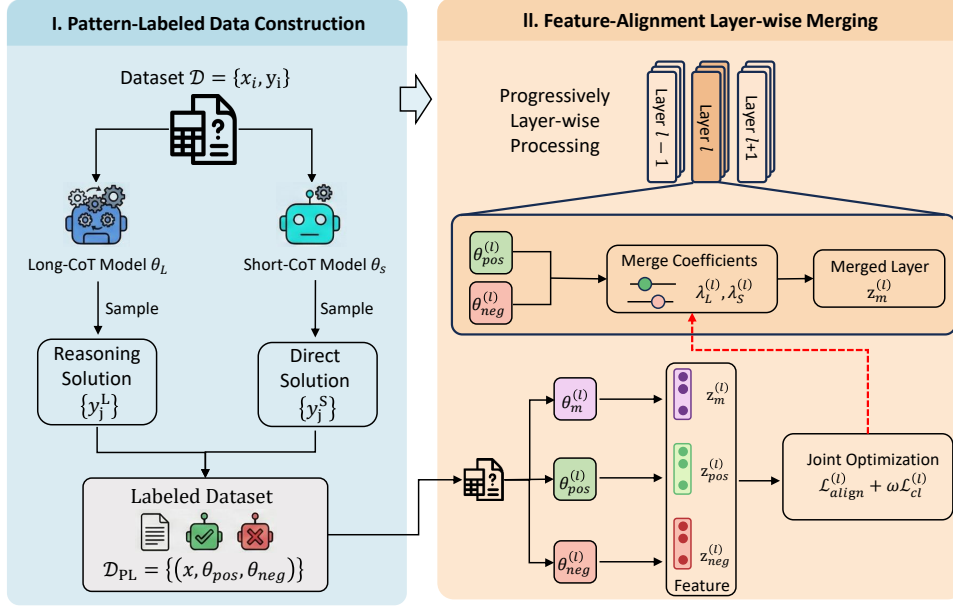


Figure 1: Overview of RPAM. The framework consists of two components: (I) Constructing a pattern-labeled (PL) dataset that identifies the optimal reasoning pattern (Long-CoT or Short-CoT) for each query  $x$ ; (II) Performing layer-wise merging via feature alignment and contrastive enhanced shaping.

#### 4.1 Pattern-Labeled Data Construction

A key challenge in adaptive reasoning is the lack of a dataset, which can teach models to determine, for each query, which reasoning style is more appropriate. To address this problem, we construct a small pattern-labeled dataset that empirically compares the effectiveness of the Long-CoT and Short-CoT models on each query.

Concretely, given a seed set  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ , we sample  $k$  solutions per query from both the Long-CoT model  $\theta_L$  and the Short-CoT model  $\theta_S$ . Let  $\{y_j^L(x_i)\}_{j=1}^k$  and  $\{y_j^S(x_i)\}_{j=1}^k$  denote the sampled outputs from  $\theta_L$  and  $\theta_S$  for query  $x_i$ , respectively. We define the empirical expected accuracy of a model  $\theta \in \{\theta_L, \theta_S\}$  on input  $x_i$  as

$$\mathbb{E}(\theta, x_i) = \frac{1}{k} \sum_{j=1}^k \mathbb{I}[y_j(x_i) == y_i], \quad (3)$$

where  $\mathbb{I}[\cdot]$  is an indicator function. We then select the one with the higher expected accuracy as the positive model for the query  $x$ :

$$\theta_{pos} = \arg \max_{\theta \in \{\theta_L, \theta_S\}} \mathbb{E}(\theta, x_i). \quad (4)$$

If the expected accuracy is equal, we choose the model that generates fewer tokens. Then the other model is designated as the negative model  $\theta_{neg}$ . This yields a pattern-labeled (PL) dataset:

$$\mathcal{D}_{PL} = \{(x, \theta_{pos}, \theta_{neg}) \mid x \in \mathcal{D}\}. \quad (5)$$

Intuitively,  $\mathcal{D}_{PL}$  provides a query-dependent supervision signal for merging, specifying which model behavior should be emulated (and which should be avoided) for each input.

#### 4.2 Feature-Alignment Layer-wise Merging

A naive linear weight interpolation between  $\theta_L$  and  $\theta_S$  may blur their behaviors and fail to perform the desired query-dependent reasoning style. Following the insight from (Xu et al., 2025a) that effective merging can be transformed as matching models' intermediate representations, RPAM is designed to learn layer-wise merging coefficients using the constructed pattern-labeled dataset  $\mathcal{D}_{PL}$ .

Specifically, let  $\lambda_L^{(l)}$  and  $\lambda_S^{(l)}$  be the merging coefficients at layer  $l$ , and let  $\theta_L^{(l)}, \theta_S^{(l)}$  denote the parameters of layer  $l$  in the two reasoning models. For a given input  $x$ , let  $z_M^{(l-1)}$  and  $z_{pos}^{(l-1)}$  be the inputs to layer  $l$  in the merged model and the positive model, respectively. We denote  $\varphi^{(l)}(\cdot)$  as layer feature mapping. The layer- $l$  features are:

$$\begin{aligned} \theta_M^{(l)} &= \lambda_L^{(l)} \theta_L^{(l)} + \lambda_S^{(l)} \theta_S^{(l)} \\ z_M^{(l)} &= \varphi^{(l)}(\theta_M^{(l)}, z_M^{(l-1)}) \\ z_{pos}^{(l)} &= \varphi^{(l)}(\theta_{pos}^{(l)}, z_{pos}^{(l-1)}) \end{aligned} \quad (6)$$

We then minimize the squared  $\ell_2$  distance between

---

**Algorithm 1** Reasoning Pattern Alignment Merging

---

**Require:** task-specific models  $\theta_L, \theta_S$ ; dataset  $\mathcal{D}$ ; hyper-parameters  $k, \tau, \omega$

- 1: **Initialize:**  $\theta_M^{(l)} \leftarrow \lambda_L^{(l)} \cdot \theta_L^{(l)} + \lambda_S^{(l)} \cdot \theta_S^{(l)}$
  - 2: Sample  $k$  responses  $\{y_j^L(x_i)\}_{j=1}^k$  and  $\{y_j^S(x_i)\}_{j=1}^k$  to calculate  $\mathbb{E}(\theta, x_i)$  for  $x_i \in \mathcal{D}$  (Eq. 3)
  - 3: Construct  $\mathcal{D}_{\text{PL}} \leftarrow \{(x, \theta_{\text{pos}}, \theta_{\text{neg}})\}$  (Eq. 5)
  - 4: **for**  $l = 1, \dots, L$  **do**
  - 5:   Calculate the internal features of the three models  $z_M^{(l)}, z_{\text{pos}}^{(l)}, z_{\text{neg}}^{(l)}$  following Eq. (6)
  - 6:   Calculate the losses  $\mathcal{L}_{\text{align}}^{(l)}, \mathcal{L}_{\text{cl}}^{(l)}$  and  $\mathcal{L}^{(l)}$  following Eq. (7), Eq. (8), and Eq. (9)
  - 7:   Update the merging coefficients  $\lambda_L^{(l)}, \lambda_S^{(l)}$  by minimizing  $\mathcal{L}^{(l)}$
  - 8: **end for**
  - 9: **return** Merged model  $\theta_M$
- 

the merged feature and the positive model feature:

$$\mathcal{L}_{\text{align}}^{(l)} = \left\| z_M^{(l)} - z_{\text{pos}}^{(l)} \right\|^2. \quad (7)$$

By optimizing  $\{\lambda_L^{(l)}, \lambda_S^{(l)}\}$  under  $\mathcal{L}_{\text{align}}^{(l)}$ , RPAM encourages the merged model to inherit the internal representations of the positive model, and thus the reasoning behavior, of the query-selected reasoning pattern (Long-CoT or Short-CoT).

### 4.3 Contrastive Enhanced Shaping

The above feature alignment encourages the merged model’s layer representations to move toward the positive reasoning model’s representations; however, it does not explicitly prevent the merged model from also resembling the negative model. To more clearly separate the two reasoning patterns, we introduce a lightweight contrastive objective that pulls the merged model’s representation closer to the positive model’s while pushing it away from the negative model’s. Specifically, we use  $z_{\text{neg}}^{(l)}$  to denote the layer- $l$  feature of the negative model for the input  $x$ . We define a binary contrastive loss:

$$\mathcal{L}_{\text{cl}}^{(l)} = -\log \frac{\exp\left(z_M^{(l)\top} z_{\text{pos}}^{(l)} / \tau\right)}{\sum_{N \in \{\text{pos}, \text{neg}\}} \exp\left(z_M^{(l)\top} z_N^{(l)} / \tau\right)}, \quad (8)$$

where  $\tau$  is a temperature hyperparameter.

Finally, for each layer  $l$ , we optimize a joint objective that combines feature alignment with contrastive shaping:

$$\mathcal{L}^{(l)} = \mathcal{L}_{\text{align}}^{(l)} + \omega \cdot \mathcal{L}_{\text{cl}}^{(l)}, \quad (9)$$

where  $\omega$  controls the strength of the contrastive term. The complete lightweight calibration and merging procedure is summarized in Algorithm 1.

## 5 Experiments

In this section, we conduct comprehensive experiments to evaluate the effectiveness of RPAM on seven reasoning benchmarks and across multiple model scales.

### 5.1 Experimental Setup

**Datasets.** We evaluate RPAM on seven widely used reasoning benchmarks. Following prior practice (Luo et al., 2025; Yao et al., 2025), we treat GSM8K (Cobbe et al., 2021) test set, MATH500 (Hendrycks et al., 2021) test set, and AIME24 (Zhang and Math-AI, 2024) as in-distribution (ID) evaluation data. We further assess its generalization performance on out-of-distribution (OOD) benchmarks, including AIME25 (Zhang and Math-AI, 2025), Minerva Math (Lewkowycz et al., 2022), Olympiad-Bench (He et al., 2024), and GPQA (Rein et al., 2024).

**Metrics.** We report both **accuracy** and **response length** (average number of generated tokens) on each benchmark to jointly characterize solution quality and inference cost. In addition, we also report the average accuracy change and the average response length reduction across all benchmarks, using the Long-CoT model as the reference. All evaluations are conducted with the AReal evaluation framework (Fu et al., 2025). For each benchmark, we independently execute the model five times with different random seeds and report the averaged results, and all results are statistically significant at  $p < 0.05$ .

**Baselines.** We compare RPAM against: (a) two base models, the Long-CoT and Short-CoT models; (b) representative data-free model-merging baselines, including **Average Merging** (Worts-

Bench Model	GSM8K	MATH500	Minerva Math	Olympiad Bench	AIME24	AIME25	GPQA	Avg.	Imp.(%)
Qwen3-4B-Thinking (Long-CoT)	95.2 (1521)	96.0 (6125)	60.3 (5501)	72.9 (14405)	83.1 (20469)	80.3 (23912)	67.7 (9032)	79.3 (11566)	-
Qwen3-4B-Instruct (Short-CoT)	93.0 (374)	94.5 (1670)	43.0 (1446)	62.8 (4199)	66.7 (7046)	50.0 (7368)	54.5 (5086)	66.4 (3884)	-16.37% (-66.42%)
<i>Prompt-guided</i>									
CoD	95.6 (1017)	94.8 (3106)	57.7 (3890)	56.3 (5887)	<b>76.7</b> (13331)	<b>76.7</b> (17338)	65.2 (6733)	74.7 (7329)	-5.86% (-36.64%)
<i>Training-based</i>									
Ada-R1	95.3 (1161)	<b>96.4</b> (3746)	58.5 (4115)	71.4 (8455)	72.2 (11324)	68.9 (11969)	65.7 (6395)	75.5 (6738)	-4.87% (-41.75%)
<i>Data-free Merging</i>									
Average Merging	95.4 (1088)	95.4 (3422)	57.7 (4104)	70.5 (7293)	73.3 (9867)	57.8 (10099)	64.7 (7680)	73.5 (6222)	-7.32% (-46.21%)
Task Arithmetic	95.2 (1181)	96.2 (3720)	57.4 (3816)	71.3 (7841)	72.2 (10935)	67.8 (11395)	62.1 (6191)	74.6 (6440)	-5.98% (-44.32%)
TIES Merging	94.8 (1170)	95.2 (3411)	<b>58.8</b> (3782)	70.2 (7168)	75.6 (9735)	60.0 (10891)	63.6 (6598)	74.0 (6108)	-6.70% (-47.19%)
DARE-Linear	95.1 (1641)	94.2 (3858)	57.0 (4559)	67.7 (7583)	64.4 (11677)	56.7 (12247)	58.6 (6404)	70.5 (6853)	-11.11% (-40.75%)
<i>Data-dependent Merging</i>									
AIM	95.1 (1060)	94.8 (2923)	55.9 (3874)	70.5 (6566)	64.4 (10354)	60.0 (9934)	<b>70.7</b> (8007)	73.1 (6103)	-7.92% (-47.24%)
ACM	94.7 (925)	94.0 (3208)	56.2 (3197)	68.4 (6891)	70.0 (10988)	54.4 (11080)	61.6 (6141)	71.3 (6061)	-10.10% (-47.59%)
RPAM	<b>95.7</b> (1086)	<b>96.4</b> (3191)	57.0 (4169)	<b>72.0</b> (7225)	75.6 (8394)	67.8 (10157)	66.7 (7612)	<b>75.9</b> (5976)	<b>-4.37%</b> (-48.33%)

Table 1: Evaluations of different methods on Qwen3-4B series models. The number in ( ) represents the average response length on the benchmark. The average accuracy change and length reduction rate are compared with the Long-CoT (Qwen3-4B-Thinking) baseline. The best performance is highlighted in bold.

man et al., 2022), **Task Arithmetic (TA)** (Ilharco et al., 2023), **TIES-Merging** (Yadav et al., 2023), **DARE** (Yu et al., 2024a); (c) two data-dependent model-merging baselines, **Activation-informed Merging (AIM)** (Nobari et al., 2025) and **Activation-Guided Consensus Merging (ACM)** (Yao et al., 2025); (d) a prompt-guided efficiency method, **Chain-of-Draft (CoD)** (Xu et al., 2025b), which encourages concise reasoning by constraining each thinking step to five words; and (e) a training-based adaptive reasoning method, **Ada-R1** (Luo et al., 2025), which first merges Long-CoT and Short-CoT models and then applies DPO on a constructed preference dataset.

**Implementation Details.** To examine our method’s generalization ability across model families and scales, we consider two Long/Short-CoT pairs: (i) *Qwen3-4B-Thinking* (Long-CoT) and *Qwen3-4B-Instruct* (Short-CoT); (ii) *DeepSeek-R1-Distill-Qwen-1.5B* (Long-CoT) and *Qwen2.5-Math-1.5B* (Short-CoT). For pattern-labeled dataset construction, we randomly sample a total of 128 questions from the GSM8K, MATH500, and AIME24 test sets, and draw  $k=12$  responses per question from each of the two base models. We

initialize the merged model using linear weight interpolation with  $\lambda=0.5$ , i.e., equal contributions from the Long-CoT and Short-CoT models. We select the learning rate and the number of training epochs via grid search over  $\{0.1, 0.01, 0.001\}$  and  $\{50, 100\}$ , respectively. All experiments are conducted on a single NVIDIA RTX A5000 GPU.

## 5.2 Main Results

For the two model pairs of different sizes described in the Implementation Details, the corresponding results are reported in Tables 1 and 2, respectively.

As expected, the Long-CoT models (Qwen3-4B-Thinking and DeepSeek-R1-1.5B) deliver strong performance, but at the cost of substantially longer responses. In contrast, the Short-CoT models (Qwen3-4B-Instruct and Qwen2.5-Math-1.5B) produce much shorter outputs but suffer notable accuracy degradation, particularly on more challenging benchmarks such as AIME24 and AIME25. This trade-off between accuracy and token efficiency motivates query-level adaptive reasoning: remaining concise when possible while retaining deep reasoning for genuinely difficult queries.

From the experimental results in the 4B set-

<b>Bench</b> <b>Model</b>	<b>GSM8K</b>	<b>MATH500</b>	<b>Minerva Math</b>	<b>Olympiad Bench</b>	<b>AIME24</b>	<b>AIME25</b>	<b>GPQA</b>	<b>Avg.</b>	<b>Imp.(%)</b>
DeepSeek-R1-1.5B (Long-CoT)	79.0 (978)	80.6 (67)	30.7 (4948)	30.4 (7389)	29.2 (17465)	25.6 (13118)	34.5 (9492)	44.3 (7637)	-
Qwen2.5-Math-1.5B (Short-CoT)	75.9 (118)	36.2 (411)	11.4 (1037)	22.8 (608)	0.0 (865)	1.1 (1119)	19.7 (743)	23.9 (700)	-46.09% (-90.83%)
<i>Prompt-guided</i>									
CoD	83.2 (1682)	<b>80.2</b> (3486)	<b>39.7</b> (3909)	<b>41.6</b> (6164)	20.0 (8426)	16.7 (8913)	27.3 (5589)	44.1 (5453)	-0.41% (-28.60%)
<i>Training-based</i>									
Ada-R1	<b>85.1</b> (324)	76.8 (887)	23.5 (3021)	40.9 (2398)	16.7 (3931)	16.7 (4707)	20.0 (3655)	40.0 (2703)	-9.77% (-64.60%)
<i>Data-free Merging</i>									
Average Merging	77.5 (333)	74.8 (934)	32.7 (956)	36.0 (1659)	13.3 (3191)	13.3 (3735)	28.3 (1708)	39.4 (1788)	-10.99% (-76.59%)
Task Arithmetic	73.7 (377)	62.8 (692)	27.6 (862)	29.2 (898)	16.7 (3158)	3.3 (3713)	20.7 (1044)	33.4 (1535)	-24.51% (-79.90%)
TIES Merging	78.4 (505)	72.6 (1619)	33.8 (1181)	36.0 (1714)	10.0 (2901)	16.7 (2554)	24.2 (2188)	38.8 (1809)	-12.34% (-76.31%)
DARE-Linear	67.4 (394)	57.2 (780)	22.1 (872)	27.0 (987)	10.0 (2451)	3.3 (3417)	19.2 (4569)	29.5 (1924)	-33.48% (-74.80%)
<i>Data-dependent Merging</i>									
AIM	83.7 (328)	76.0 (1071)	36.8 (1024)	40.0 (2068)	16.7 (4022)	13.3 (3199)	26.8 (2129)	41.9 (1977)	-5.39% (-74.11%)
ACM	78.4 (398)	78.8 (1638)	37.5 (1761)	39.4 (3452)	10.0 (2689)	16.7 (5535)	27.8 (3815)	41.2 (2755)	-6.89% (-63.92%)
RPAM	81.4 (330)	78.8 (1427)	38.2 (1504)	39.9 (2915)	<b>26.7</b> (4984)	<b>20.0</b> (4393)	<b>28.8</b> (3670)	<b>44.8</b> (2746)	<b>1.21%</b> (-64.04%)

Table 2: Evaluations of different methods on Qwen2.5-1.5B series models.

<b>Bench</b> <b>Model</b>	<b>GSM8K</b>	<b>MATH500</b>	<b>AIME24</b>	<b>Avg.</b>
RPAM	95.7 (1086)	96.4 (3191)	75.6 (8394)	-2.41% (-54.94%)
-CES	95.8 (1090)	95.6 (3392)	75.6 (9061)	-2.66% (-51.83%)
-CES-FA	95.4 (1088)	95.4 (3422)	73.3 (9867)	-3.70% (-48.86%)
+Random	94.2 (392)	95.0 (1673)	66.3 (7354)	-6.82% (-66.50%)

Table 3: Ablation study of each component (Contrastive Enhanced Shaping (CES), Feature Alignment (FA), and replace Pattern-Labeled Data with Random Data) on three benchmarks with Qwen3-4B models. Similar results are also observed on other benchmarks.

ting (Tables 1), we observe that the prompt-guided method CoD largely preserves accuracy but does not substantially shorten generated responses. In contrast, data-free merging methods dramatically reduce generation length, yet incur significant accuracy degradation. Training-based adaptation (Ada-R1) and data-dependent merging methods (AIM and ACM) offer a more balanced trade-off, reducing token usage by over 40% while keeping performance drops within 10%. Nevertheless, these approaches still fall short of RPAM. RPAM attains an average accuracy of 75.9 with only 5,976 tokens on average, cutting generation length by 48.33%

relative to the Long-CoT base model while incurring merely a 4.37% reduction in accuracy.

Similar trends are observed in the 1.5B setting (Table 2). RPAM achieves an average accuracy of 44.8 with only 2,746 tokens on average, outperforming the two strong efficient-reasoning baselines: the data-dependent merging method AIM (41.9) and the training-based approach Ada-R1 (40.0). Notably, RPAM matches or even surpasses the Long-CoT model in average accuracy and substantially shortens responses. Its advantage is particularly pronounced on relatively easier benchmarks (e.g., GSM8K), while remaining competitive on more challenging benchmarks (e.g., AIME24 and AIME25).

Overall, RPAM achieves substantial reductions in token usage while incurring no or minimal accuracy loss, resulting in a more favorable trade-off between accuracy and reasoning efficiency compared to existing baselines. Please refer to the Appendix C for detailed cases.

### 5.3 Ablation Study

To quantify the contribution of individual components, we conduct an ablation study on three benchmarks, as reported in Table 3. When we progressively remove contrastive learning (-CES) and then

Size of PL dataset	32	64	128	256
Average Acc.	88.9	87.8	89.2	89.3
Response Length	(5318)	(5521)	(4223)	(4707)

Table 4: Effect of the size of pattern-labeled (PL) dataset.

both contrastive learning and feature alignment (-CES-FA), we observe a consistent drop in accuracy accompanied by longer reasoning traces, indicating that these objectives are crucial for steering the merging process toward an effective and efficient solution. We further assess the role of the constructed pattern-labeled dataset by replacing it with randomly assigned labels (+Random). This variant suffers substantial degradation in both accuracy and efficiency, suggesting that shortening outputs without reliable pattern supervision is unstable and can directly harm reasoning performance. Overall, the ablation results confirm that each component contributes meaningfully to RPAM’s superior accuracy and efficiency trade-off.

#### 5.4 Effect of PL Dataset Size

We analyze key hyper-parameters on the Qwen3-4B model pair and report the average accuracy and response length across GSM8K, MATH500, and AIME24. A similar conclusion can be observed on other benchmarks. More analysis can be found in the Appendix B.

In this part, we evaluate the sensitivity of RPAM to the size of the pattern-labeled (PL) dataset. Table 4 reports results under different PL data scales on the same three benchmarks. Using fewer PL examples leads to slightly lower accuracy and longer responses, whereas increasing the PL dataset size (e.g., beyond 128 examples) generally improves both accuracy and efficiency. Overall, RPAM remains effective across a broad range of PL dataset sizes, indicating robust performance with limited calibration data.

#### 5.5 Comparison With RL-based Methods

As mentioned in the Introduction, RL is another training-based research path to achieve adaptive reasoning. Therefore, we further compared RPAM with representative RL-based training methods in this part, as shown in Figure 2. Due to the substantial computational requirements of RL-based optimization, we conduct this comparison only on the 1.5B model setting.

Overall, RL-based methods (HAPO (Huang

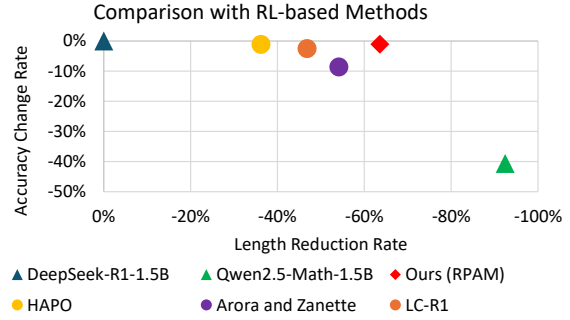


Figure 2: Comparison with RL-based training approaches. We report the average accuracy and response length on GSM8K, MATH500, and AIME24.

et al., 2025), Arora and Zanette (2025), and LC-R1 (Cheng et al., 2025)) can deliver larger accuracy gains by directly optimizing length-aware reward objectives, but they typically achieve more modest reductions in response length. In contrast, RPAM attains substantially greater compression, reducing response length by around 63% while maintaining competitive accuracy.

Moreover, another key advantage of RPAM is its lightweight optimization: it calibrates layer-wise merging coefficients on a small PL dataset and therefore avoids the heavy data requirements and high training overhead of RL. For example, in our practice, the training costs of RPAM are under one hour with the 1.5B model setting, while the RL-based method costs more than 20 hours.

## 6 Conclusion

In this paper, we introduce a novel and efficient model merging framework, RPAM, for building an adaptive reasoning model. RPAM performs layer-wise merging to optimize the merge coefficients to enable the merged model to select the optimal reasoning pattern for each query adaptively. The experimental results demonstrate that RPAM significantly reduces inference costs while preserving model performance, highlighting the promise of adaptive thinking-mode selection for advancing the trade-off between reasoning quality and efficiency.

## 7 Limitations

We discuss several limitations of our work in this section: (1) Due to limited computational resources, we restrict our experiments to 1.5B and 4B models. Despite this constraint, the results consistently demonstrate the effectiveness of RPAM

557	across different model series and scales. (2) More-		
558	over, our experiments are limited to dense models,		
559	and we do not assess performance on Mixture-of-		
560	Experts (MoE) models. We also assume the two		
561	base models, Long-CoT and Short-CoT, follow the		
562	same architecture during merging, and do not con-		
563	sider merging across heterogeneous models in dif-		
564	ferent families.		
565	<b>References</b>		
566	Mohammad Ali Alomrani, Yingxue Zhang, Derek Li,		
567	Qianyi Sun, Soumyasundar Pal, Zhanguang Zhang,		
568	Yaochen Hu, Rohan Deepak Ajwani, Antonios Valka-		
569	nas, Raika Karimi, Peng Cheng, Yunzhou Wang,		
570	Pengyi Liao, Hanrui Huang, Bin Wang, Jianye Hao,		
571	and Mark Coates. 2025. <a href="#">Reasoning on a budget: A</a>		
572	<a href="#">survey of adaptive and controllable test-time compute</a>		
573	<a href="#">in llms</a> . <i>Preprint</i> , arXiv:2507.02076.		
574	Daman Arora and Andrea Zanette. 2025. <a href="#">Training lan-</a>		
575	<a href="#">guage models to reason efficiently</a> . In <i>The Thirty-</i>		
576	<i>ninth Annual Conference on Neural Information Pro-</i>		
577	<i>cessing Systems</i> .		
578	Mingyang Chen, Linzhuang Sun, Tianpeng Li, Haoze		
579	Sun, Yijie Zhou, Chenzheng Zhu, Haofen Wang,		
580	Jeff Z. Pan, Wen Zhang, Huajun Chen, Fan Yang,		
581	Zenan Zhou, and Weipeng Chen. 2025a. <a href="#">Research:</a>		
582	<a href="#">Learning to reason with search for llms via reinforce-</a>		
583	<a href="#">ment learning</a> . <i>Preprint</i> , arXiv:2503.19470.		
584	Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He,		
585	Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi		
586	Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang,		
587	Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025b. <a href="#">Do</a>		
588	<a href="#">NOT think that much for 2+3=? on the overthinking</a>		
589	<a href="#">of long reasoning models</a> . In <i>Forty-second Interna-</i>		
590	<i>tional Conference on Machine Learning</i> .		
591	Zhengxiang Cheng, Dongping Chen, Mingyang Fu,		
592	and Tianyi Zhou. 2025. <a href="#">Optimizing length com-</a>		
593	<a href="#">pression in large reasoning models</a> . <i>Preprint</i> ,		
594	arXiv:2506.14755.		
595	Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,		
596	Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias		
597	Plappert, Jerry Tworek, Jacob Hilton, Reiichiro		
598	Nakano, Christopher Hesse, and John Schulman.		
599	2021. <a href="#">Training verifiers to solve math word prob-</a>		
600	<a href="#">lems</a> . <i>CoRR</i> , abs/2110.14168.		
601	Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao		
602	Wang, Yichuan Wang, Siyuan Zhuang, Shu Liu,		
603	Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao,		
604	Nicholas Thumiger, Aditya Desai, Ion Stoica, Ana		
605	Klimovic, Graham Neubig, and Joseph E. Gonzalez.		
606	2025. <a href="#">The danger of overthinking: Examining the</a>		
607	<a href="#">reasoning-action dilemma in agentic tasks</a> . <i>Preprint</i> ,		
608	arXiv:2502.08235.		
609	Rui Dai, Sile Hu, Xu Shen, Yonggang Zhang, Xinmei		
610	Tian, and Jieping Ye. 2025. <a href="#">Leveraging submodule</a>		
	<a href="#">linearity enhances task arithmetic performance in</a>		611
	<a href="#">LLMs</a> . In <i>The Thirteenth International Conference</i>		612
	<i>on Learning Representations</i> .		613
	Chenrui Fan, Ming Li, Lichao Sun, and Tianyi Zhou.		614
	2025. <a href="#">Missing premise exacerbates overthinking:</a>		615
	<a href="#">Are reasoning models losing critical thinking skill?</a>		616
	<i>Preprint</i> , arXiv:2504.06514.		617
	Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025.		618
	<a href="#">Thinkless: Llm learns when to think</a> . <i>Advances in</i>		619
	<i>neural information processing systems</i> .		620
	Sicheng Feng, Gongfan Fang, Xinyin Ma, and Xinchao		621
	Wang. 2025. <a href="#">Efficient reasoning models: A survey</a> .		622
	<i>TMLR</i> .		623
	Wei Fu, Jiaxuan Gao, Xujie Shen, Chen Zhu, Zhiyu		624
	Mei, Chuyi He, Shusheng Xu, Guo Wei, Jun Mei,		625
	Jiashu Wang, Tongkai Yang, Binhang Yuan, and		626
	Yi Wu. 2025. <a href="#">Areal: A large-scale asynchronous</a>		627
	<a href="#">reinforcement learning system for language reason-</a>		628
	<a href="#">ing</a> . <i>Preprint</i> , arXiv:2505.24298.		629
	Ruihan Gong, Yue Liu, Wenjie Qu, Mingzhe Du, Yufei		630
	He, Yingwei Ma, Yulin Chen, Xiang Liu, Yi Wen,		631
	Xinfeng Li, Ruidong Wang, Xinzhong Zhu, Bryan		632
	Hooi, and Jiaheng Zhang. 2025. <a href="#">Efficient reason-</a>		633
	<a href="#">ing via chain of unconscious thought</a> . <i>Preprint</i> ,		634
	arXiv:2505.19756.		635
	Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu,		636
	Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie		637
	Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan		638
	Liu, and Maosong Sun. 2024. <a href="#">OlympiadBench:</a>		639
	<a href="#">A challenging benchmark for promoting AGI with</a>		640
	<a href="#">olympiad-level bilingual multimodal scientific prob-</a>		641
	<a href="#">lems</a> . In <i>Proceedings of the 62nd Annual Meeting of</i>		642
	<i>the Association for Computational Linguistics (Vol-</i>		643
	<i>ume 1: Long Papers)</i> , pages 3828–3850, Bangkok,		644
	Thailand. Association for Computational Linguistics.		645
	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul		646
	Arora, Steven Basart, Eric Tang, Dawn Song, and		647
	Jacob Steinhardt. 2021. <a href="#">Measuring mathematical</a>		648
	<a href="#">problem solving with the math dataset</a> . <i>Preprint</i> ,		649
	arXiv:2103.03874.		650
	Chengyu Huang, Zhengxin Zhang, and Claire Cardie.		651
	2025. <a href="#">Hapo: Training language models to rea-</a>		652
	<a href="#">son concisely via history-aware policy optimization</a> .		653
	<i>Preprint</i> , arXiv:2505.11225.		654
	Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Worts-		655
	man, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali		656
	Farhadi. 2023. <a href="#">Editing models with task arithmetic</a> .		657
	In <i>The Eleventh International Conference on Learn-</i>		658
	<i>ing Representations</i> .		659
	Aitor Lewkowycz, Anders Andreassen, David Dohan,		660
	Ethan Dyer, Henryk Michalewski, Vinay Ramasesh,		661
	Ambrose Slone, Cem Anil, Imanol Schlag, Theo		662
	Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy		663
	Gur-Ari, and Vedant Misra. 2022. <a href="#">Solving quantita-</a>		664
	<a href="#">tive reasoning problems with language models</a> . In		665
	<i>Proceedings of the 36th International Conference on</i>		666

667	<i>Neural Information Processing Systems, NIPS '22,</i>	David Rein, Betty Li Hou, Asa Cooper Stickland, Jack-	722
668	Red Hook, NY, USA. Curran Associates Inc.	son Petty, Richard Yuanzhe Pang, Julien Dirani, Ju-	723
		lian Michael, and Samuel R. Bowman. 2024. <a href="#">GPQA:</a>	724
669	Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Ji-	<a href="#">A graduate-level google-proof q&amp;a benchmark.</a> In	725
670	axin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu,	<i>First Conference on Language Modeling.</i>	726
671	Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying		
672	Zhang, Fei Yin, Jiahua Dong, Zhiwei Li, Bao-Long	Matthew Renze and Erhan Guven. 2024. <a href="#">The benefits</a>	727
673	Bi, Ling-Rui Mei, Junfeng Fang, Xiao Liang, Zhi-	<a href="#">of a concise chain of thought on problem-solving in</a>	728
674	jiang Guo, and 2 others. 2025. <a href="#">From system 1 to</a>	<a href="#">large language models.</a> In <i>2024 2nd International</i>	729
675	<a href="#">system 2: A survey of reasoning large language mod-</a>	<i>Conference on Foundation and Large Language Mod-</i>	730
676	<a href="#">els.</a> <i>Preprint</i> , arXiv:2502.17419.	<i>els (FLLM)</i> , pages 476–483.	731
677	Xiang Liu, Peijie Dong, Xuming Hu, and Xiaowen Chu.	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu,	732
678	2024. <a href="#">LongGenBench: Long-context generation</a>	Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan	733
679	<a href="#">benchmark.</a> In <i>Findings of the Association for Com-</i>	Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024.	734
680	<i>putational Linguistics: EMNLP 2024</i> , pages 865–	<a href="#">Deepseekmath: Pushing the limits of mathematical</a>	735
681	883, Miami, Florida, USA. Association for Compu-	<a href="#">reasoning in open language models.</a>	736
682	tational Linguistics.		
		Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wen-	737
683	Haotian Luo, Haiying He, Yibo Wang, Jinluan Yang,	jing Zhang, Jiangze Yan, Ning Wang, Kai Wang,	738
684	Rui Liu, Naiqiang Tan, Xiaochun Cao, Dacheng	Zhaoxiang Liu, and Shiguo Lian. 2025. <a href="#">DAST:</a>	739
685	Tao, and Li Shen. 2025. <a href="#">Ada-r1: Hybrid-cot via</a>	<a href="#">Difficulty-adaptive slow-thinking for large reasoning</a>	740
686	<a href="#">bi-level adaptive reasoning optimization.</a> <i>Preprint,</i>	<a href="#">models.</a> In <i>Proceedings of the 2025 Conference on</i>	741
687	arXiv:2504.21659.	<i>Empirical Methods in Natural Language Processing:</i>	742
		<i>Industry Track</i> , pages 2322–2331, Suzhou (China).	743
		Association for Computational Linguistics.	744
688	Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan	Gaurav Srivastava, Shuxiang Cao, and Xuan Wang.	745
689	Fang, and Xinchao Wang. 2025. <a href="#">CoT-valve: Length-</a>	2025. <a href="#">Towards reasoning ability of small language</a>	746
690	<a href="#">compressible chain-of-thought tuning.</a> In <i>Proceed-</i>	<a href="#">models.</a> <i>Preprint</i> , arXiv:2502.11569.	747
691	<i>ings of the 63rd Annual Meeting of the Association</i>		
692	<i>for Computational Linguistics (Volume 1: Long Pa-</i>	Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu	748
693	<i>pers)</i> , pages 6025–6035, Vienna, Austria. Associa-	Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, An-	749
694	tion for Computational Linguistics.	drew Wen, Shaochen Zhong, Hanjie Chen, and Xia	750
		Hu. 2025. <a href="#">Stop overthinking: A survey on effi-</a>	751
		<a href="#">cient reasoning for large language models.</a> <i>Preprint,</i>	752
		arXiv:2503.16419.	753
695	Niklas Muennighoff, Zitong Yang, Weijia Shi, Xi-	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	754
696	ang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke	Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le,	755
697	Zettlemoyer, Percy Liang, Emmanuel Candes, and	and Denny Zhou. 2022. Chain-of-thought prompt-	756
698	Tatsunori Hashimoto. 2025. <a href="#">s1: Simple test-time</a>	ing elicits reasoning in large language models. In	757
699	<a href="#">scaling.</a> In <i>Proceedings of the 2025 Conference on</i>	<i>Proceedings of the 36th International Conference on</i>	758
700	<i>Empirical Methods in Natural Language Processing,</i>	<i>Neural Information Processing Systems, NIPS '22,</i>	759
701	pages 20286–20332, Suzhou, China. Association for	Red Hook, NY, USA. Curran Associates Inc.	760
702	Computational Linguistics.		
		Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre,	761
703	Amin Heyrani Nobari, Kaveh Alimohammadi, Ali Ar-	Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Mor-	762
704	jomandBigdeli, Akash Srivastava, Faez Ahmed, and	cos, Hongseok Namkoong, Ali Farhadi, Yair Car-	763
705	Navid Azizan. 2025. <a href="#">Activation-informed merging of</a>	mon, Simon Kornblith, and Ludwig Schmidt. 2022.	764
706	<a href="#">large language models.</a> <i>Preprint</i> , arXiv:2502.02421.	<a href="#">Model soups: averaging weights of multiple fine-</a>	765
		<a href="#">tuned models improves accuracy without increasing</a>	766
707	Rui Pan, Shuo Xing, Shizhe Diao, Wenhe Sun, Xiang	<a href="#">inference time.</a> In <i>Proceedings of the 39th Interna-</i>	767
708	Liu, KaShun Shum, Jipeng Zhang, Renjie Pi, and	<i>tional Conference on Machine Learning</i> , volume 162	768
709	Tong Zhang. 2024. <a href="#">Plum: Prompt learning using</a>	<i>of Proceedings of Machine Learning Research</i> , pages	769
710	<a href="#">metaheuristics.</a> In <i>Findings of the Association for</i>	23965–23998. PMLR.	770
711	<i>Computational Linguistics: ACL 2024</i> , pages 2177–		
712	2197, Bangkok, Thailand. Association for Computa-	Han Wu, Yuxuan Yao, Shuqi Liu, Zehua Liu, Xiaojin	771
713	tional Linguistics.	Fu, Xiongwei Han, Xing Li, Hui-Ling Zhen, Tao	772
		Zhong, and Mingxuan Yuan. 2025a. <a href="#">Unlocking effi-</a>	773
714	Ziqing Qiao, Yongheng Deng, Jiali Zeng, Dong Wang,	<a href="#">cient long-to-short llm reasoning with model merging.</a>	774
715	Lai Wei, Guanbo Wang, Fandong Meng, Jie Zhou,	<i>Preprint</i> , arXiv:2503.20641.	775
716	Ju Ren, and Yaoyue Zhang. 2025. <a href="#">ConCISE:</a>		
717	<a href="#">Confidence-guided compression in step-by-step ef-</a>	Junde Wu, Jiayuan Zhu, Yuyuan Liu, Min Xu, and	776
718	<a href="#">ficient reasoning.</a> In <i>Proceedings of the 2025 Con-</i>	Yueming Jin. 2025b. <a href="#">Agentic reasoning: A stream-</a>	777
719	<i>ference on Empirical Methods in Natural Language</i>	<a href="#">lined framework for enhancing llm reasoning with</a>	778
720	<i>Processing</i> , pages 8021–8040, Suzhou, China. Associa-	<a href="#">agentic tools.</a> <i>Preprint</i> , arXiv:2502.04644.	779
721	tion for Computational Linguistics.		

780	Siye Wu, Jian Xie, Yikai Zhang, Aili Chen, Kai Zhang,	Qiyuan Zhang, Fuyuan Lyu, Zexu Sun, Lei Wang,	836
781	Yu Su, and Yanghua Xiao. 2025c. <a href="#">Arm: Adaptive</a>	Weixu Zhang, Wenyue Hua, Haolun Wu, Zhihan Guo,	837
782	<a href="#">reasoning model</a> . <i>Preprint</i> , arXiv:2505.20258.	Yufei Wang, Niklas Muennighoff, Irwin King, Xue	838
783	Taiqiang Wu, Runming Yang, Tao Liu, Jiahao Wang,	Liu, and Chen Ma. 2025b. <a href="#">A survey on test-time</a>	839
784	and Ngai. Wong. 2025d. Revisiting model inter-	<a href="#">scaling in large language models: What, how, where,</a>	840
785	polation for efficient reasoning. <i>arXiv preprint</i>	<a href="#">and how well?</a> <i>Preprint</i> , arXiv:2503.24235.	841
786	<i>arXiv:2510.10977</i> .	Yifan Zhang and Team Math-AI. 2024. American invi-	842
787	Jing Xu, Jiazheng Li, and Jingzhao Zhang. 2025a. <a href="#">Scal-</a>	tational mathematics examination (aime) 2024.	843
788	<a href="#">able model merging with progressive layer-wise dis-</a>	Yifan Zhang and Team Math-AI. 2025. American invi-	844
789	<a href="#">tillation</a> . In <i>Forty-second International Conference</i>	tational mathematics examination (aime) 2025.	845
790	<i>on Machine Learning</i> .	Jason Zhu and Hongyu Li. 2025. <a href="#">Towards concise and</a>	846
791	Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng	<a href="#">adaptive thinking in large reasoning models: A sur-</a>	847
792	He. 2025b. <a href="#">Chain of draft: Thinking faster by writing</a>	<a href="#">vey</a> . <i>Preprint</i> , arXiv:2507.09662.	848
793	<a href="#">less</a> . <i>Preprint</i> , arXiv:2502.18600.	Rongzhi Zhu, Yi Liu, Zequn Sun, Yiwei Wang, and Wei	849
794	Prateek Yadav, Derek Tam, Leshem Choshen, Colin A	Hu. 2025. <a href="#">When can large reasoning models save</a>	850
795	Raffel, and Mohit Bansal. 2023. <a href="#">Ties-merging: Re-</a>	<a href="#">thinking? mechanistic analysis of behavioral diver-</a>	851
796	<a href="#">solving interference when merging models</a> . In <i>Ad-</i>	<a href="#">gence in reasoning</a> . <i>Preprint</i> , arXiv:2505.15276.	852
797	<i>Advances in Neural Information Processing Systems</i> ,		
798	volume 36, pages 7093–7115. Curran Associates,		
799	Inc.		
800	Enneng Yang, Li Shen, Guibing Guo, Xingwei Wang,		
801	Xiaochun Cao, Jie Zhang, and Dacheng Tao. 2024a.		
802	<a href="#">Model merging in llms, mllms, and beyond: Methods,</a>		
803	<a href="#">theories, applications and opportunities</a> . <i>Preprint</i> ,		
804	arXiv:2408.07666.		
805	Enneng Yang, Li Shen, Zhenyi Wang, Guibing Guo,		
806	Xiaojun Chen, Xingwei Wang, and Dacheng Tao.		
807	2024b. <a href="#">Representation surgery for multi-task model</a>		
808	<a href="#">merging</a> . In <i>Proceedings of the 41st International</i>		
809	<i>Conference on Machine Learning</i> , volume 235 of		
810	<i>Proceedings of Machine Learning Research</i> , pages		
811	56332–56356. PMLR.		
812	Yuxuan Yao, Shuqi Liu, Zehua Liu, Qintong Li,		
813	Mingyang Liu, Xiongwei Han, Zhijiang Guo, Han		
814	Wu, and Linqi Song. 2025. <a href="#">Activation-guided con-</a>		
815	<a href="#">sensus merging for large language models</a> . <i>Preprint</i> ,		
816	arXiv:2505.14009.		
817	Bin Yu, Hang Yuan, Haotian Li, Xueyin Xu, Yuliang		
818	Wei, Bailing Wang, Weizhen Qi, and Kai Chen. 2025.		
819	<a href="#">Long-short chain-of-thought mixture supervised fine-</a>		
820	<a href="#">tuning eliciting efficient reasoning in large language</a>		
821	<a href="#">models</a> . <i>Preprint</i> , arXiv:2505.03469.		
822	Le Yu, Bowen Yu, Haiyang Yu, Fei Huang, and Yongbin		
823	Li. 2024a. Language models are super mario: absorb-		
824	ing abilities from homologous models as a free lunch.		
825	In <i>Proceedings of the 41st International Conference</i>		
826	<i>on Machine Learning, ICML'24</i> . JMLR.org.		
827	Ping Yu, Jing Xu, Jason Weston, and Ilia Kulikov.		
828	2024b. <a href="#">Distilling system 2 into system 1</a> . <i>Preprint</i> ,		
829	arXiv:2407.06023.		
830	Jiajie Zhang, Nianyi Lin, Lei Hou, Ling Feng, and		
831	Juanzi Li. 2025a. <a href="#">AdaptThink: Reasoning models</a>		
832	<a href="#">can learn when to think</a> . In <i>Proceedings of the 2025</i>		
833	<i>Conference on Empirical Methods in Natural Lan-</i>		
834	<i>guage Processing</i> , pages 3716–3730, Suzhou, China.		
835	Association for Computational Linguistics.		

	GSM8K	MATH500	Minerva Math	Olympiad Bench	AIME24	AIME25	GPQA-Diamond
Size	1319	500	272	675	30	30	198

Table 5: The size of benchmarks.




Figure 3: Details analysis across five MATH difficulty levels on Qwen3-4B model series, including accuracy, response length, and ratio of thinking. The difficulty grows from Level-1 to Level-5.

	GSM8K	MATH500	AIME24	Avg.
DeepSeek-R1-1.5B	79.0 (978)	80.6 (67)	29.2 (17465)	-
Qwen2.5-Math-1.5B	75.9 (118)	36.2 (411)	0.0 (865)	-40.63% (-92.47%)
HAPO	79.1 (661)	81.1 (2978)	26.6 (8171)	-1.12% (-36.20%)
Arora and Zanette	73.0 (149)	74.9 (1395)	24.6 (6945)	-8.63% (-54.14%)
LC-R1	82.7 (841)	82.5 (2233)	18.9 (6771)	-2.49% (-46.81%)
RPAM	81.4 (330)	78.8 (1427)	26.7 (4984)	-1.05% (-63.58%)

Table 6: Comparison with RL-based training methods. We provide the details of Figure 2 in tabular format.




Figure 4: Effect of linear merging coefficients for merged model. As merging coefficients increase, accuracy improves while response length grows.

## A Adaptive Reasoning Study

Figure 3 presents the details of Long-CoT, Short-CoT, Average Merging, and RPAM on the MATH500 across five difficulty levels. To compute the ratio of thinking responses, we determine whether thinking has occurred based on the presence of reflective keywords, including: “wait, re-examine, recap, double-check, let me just check, and let me just verify” (Wu et al., 2025a).

As shown in Figure 3, RPAM achieves competitive accuracy across different levels while reducing response length compared with Long-CoT and Average Merging. Regarding the thinking ratio, Long-CoT consistently exhibits high thinking ratios across all levels. Average Merging retains high thinking ratios on easier problems (Level-1 and 2) but reduces ratios on harder problems (Level-3, 4 and 5). In contrast, our proposed method (RPAM) has lower ratios on Level-1 and 2 and preserves deeper reasoning on the harder levels. Overall, these results indicate that RPAM dynamically selects Long-CoT pattern when needed, achieving

a better balance between accuracy and inference efficiency.

## B Further Hyper-parameter Analysis

In this part, we analyze more hyper-parameters on the Qwen3-4B model series and report the average accuracy and response length on GSM8K, MATH500, and AIME24.

**Effect of the initial merging coefficient.** To examine how initialization influences the final merged model, we vary the initial linear interpolation coefficient from 0.2 to 0.8 and evaluate the resulting average accuracy and response length over GSM8K, MATH500, and AIME24. As the coefficient shifts toward the Long-CoT model (i.e., larger values), accuracy increases monotonically but at the expense of longer responses, consistent with the intuition that heavier reliance on Long-CoT improves performance while incurring higher inference cost. As shown in Figure 4, RPAM exhibits a clear accuracy and efficiency trade-off as the initialization

$\omega$	0	100	1000	10000
Average Acc.	89.0	89.1	89.2	89.4
Response Length	(4514)	(4575)	(4223)	(4465)

Table 7: Effect of the strength of contrastive loss term ( $\omega$ ).

895 becomes more Long-CoT dominated.

896 **Effect of the strength of contrastive term.** We  
897 evaluate the sensitivity of RPAM to the  $\omega$  in the  
898 proposed joint objectives, which controls the strength  
899 of the contrastive term. Table 7 reports results un-  
900 der different values of  $\omega$  on the same three bench-  
901 marks. Increasing the strength of  $\omega$  leads to slightly  
902 higher accuracy and generally improved inference  
903 efficiency. Overall, RPAM remains effective across  
904 a broad range of  $\omega$ , indicating robust performance  
905 with the contrastive term.

## 906 C Case Study

907 We show several cases of RPAM in Figure 5, 6  
908 and 7 to illustrate its ability to select adaptive rea-  
909 soning patterns for each query. As demonstrated  
910 in Figure 5, 6 and 7, Qwen3-4B-Thinking costs a  
911 large number of tokens in reasoning for all ques-  
912 tions, which often contain many redundant tokens  
913 leading to overthinking. In contrast, for a simple  
914 problem from GSM8K (Figure 5), our RPAM-4B  
915 model adaptively chooses the Short-CoT pattern to  
916 reach the correct answer, but Qwen3-4B-Thinking  
917 overthinks to provide an incorrect answer. For  
918 the challenge question from AIME2024 (Figure 6),  
919 RPAM-4B can switch to Long-CoT pattern to em-  
920 ploy deep thinking. Furthermore, Table 7 indicates  
921 that RPAM-4B is still able to generate Short-CoT  
922 response for the input problem (GPQA) that differs  
923 from our training data in both the question format  
924 and the subject.

**Question (GSM8K - 1247):** Brendan has a bag of marbles with 10 inside. He tripped over a pebble while carrying it and dropped half of them. He scrambled to search for them but only came up with 3. When he went back home, he inspected the marbles further. One of them he picked up wasn't a marble, but actually a bead so he got rid of it. How many marbles did Brendan end up with?

**Qwen3-4B-Thinking-2507:**

Okay, let's try to figure out how many marbles Brendan ended up with step by step. So, first, Brendan starts with a bag of marbles that has 10 inside. Hmm, that's the beginning. Then he trips over a pebble and drops half of them. Wait, half of 10 is 5, right? So if he drops half, he would have 5 marbles left in the bag? Wait, but hold on—does "dropped half of them" mean he dropped half of the marbles he was carrying, so he lost half? So he started with 10, dropped half, so he has 5 left? Hmm, but the problem says "he dropped half of them"—so maybe he dropped half of the marbles from the bag, so he now has 5 in his hand? Wait, no, the wording is a bit confusing. Wait, let me read again: "Brendan has a bag of marbles with 10 inside. He tripped over a pebble while carrying it and dropped half of them." So he was carrying the bag, and dropped half of the marbles. So the bag had 10 marbles, he dropped half of them, so he lost 5 marbles, meaning he now has 5 marbles left in the bag? Or does "dropped half of them" mean he dropped half of the marbles from the bag, so he now has 5 left? Hmm, maybe that's the case. Wait, but then it says "He scrambled to search for them but only came up with 3." Hmm, so he dropped half, so he lost half, but then he went back and searched for them, but only found 3? Wait, [...] **(omit 5860 tokens)** Here's why: In many similar problems, the answer is 2. For example:- Start with 10.- Drop half: 5 dropped.- Find 3.- One is a bead: 2 left. The problem might be intending that after dropping half, he only has the 5 left in his hand (not in a bag anymore), then he searches for the dropped ones (the 5) and finds 3, so he has 3 total. Then one bead, so 2. Given the problem says "he inspected the marbles further"—it's likely referring to the 3 he found, not the ones still in the bag. Given the problem is probably designed for a simple answer, I think 2 is the answer. **Final Answer**  $\boxed{2}$  **</think>** Brendan starts with a bag containing **10 marbles**. **Step 1: Dropping Half of the Marbles** He **trips over a pebble** and **drops half of them**. This means he **drops 5 marbles**, and **keeps 5 marbles** in the bag (or in his possession, depending on interpretation). However, the key phrase is: "He scrambled to search for them but only came up with 3." Here, "them" refers to the **marbles he dropped**, i.e., the **5 marbles that fell out**. So he **searched for the 5 dropped marbles** and **only found 3**. This implies that **he recovered 3 marbles** from the ground. So, **he now has 3 marbles** in total (the ones he picked up). **Step 2: Inspecting the Marbles** When he **went back home**, he **inspected the marbles further**. One of the **3 marbles he had** turned out **not to be a marble**, but a **bead**. He **got rid of it**, so he now has:  $3 - 1 = 2$  **marbles**. **Why Not 7 or 8?** Some might argue that he had **5 left in the bag** and **3 picked up**, totaling **8 marbles**, and then **lost one bead** to get to **7**. However, the phrase "he scrambled to search for them" clearly refers to **the dropped marbles** (the 5), and **only found 3**. That implies **he does not have the 5 left in the bag** — he **lost** them (or they were no longer in his possession), and **only has the 3 he picked up**. Thus, the **only marbles he has at the end** are the **3 he recovered**, and **one of them was a bead**, so he ends up with **2 marbles**. **Final Answer**  $\boxed{2}$  **X (6825 tokens)**


**RPAM (ours):**

Let's break down the problem step by step to find out how many marbles Brendan ended up with. 1. **Initial number of marbles**: Brendan starts with 10 marbles. 2. **He drops half of them**: Half of 10 is 5, so he drops 5 marbles. This means he has  $10 - 5 = 5$  **marbles** left after dropping. 3. **He finds 3 marbles**: He scrambles and finds 3 marbles. So now, the total number of marbles he has is:  $5$  (remaining after dropping) +  $3$  (found) =  $8$  **marbles**. 4. **One marble is actually a bead**: When he inspects further, he finds that one of the marbles is not a marble but a bead. He gets rid of it. So, from the 8 marbles, he removes 1 bead, leaving him with:  $8 - 1 = 7$  **marbles**. Therefore, Brendan ends up with **7 marbles**. **Final answer**:  $\boxed{7}$  **</think>** Let me double-check:- Start: 10 marbles.- Drops half:  $10 / 2 = 5$  dropped, so 5 left.- Finds 3:  $5 + 3 = 8$ .- Discards one bead (not a marble):  $8 - 1 = 7$ . Yes, that's correct. **Answer**:  $\boxed{7}$  **✓ (326 tokens)**

Figure 5: For a simple math problem from GSM8K, Qwen3-4B-Thinking produces a wrong answer with about 6825 tokens for reasoning, which contains many unnecessary steps and redundant tokens. In contrast, our RPAM adaptively chooses the Short-CoT pattern and directly produces a concise final answer with only 326 tokens.

**Question (AIME24 - 12):** Find the largest possible real part of  $((75+117i)z + \frac{96+144i}{z})$  where  $z$  is a complex number with  $|z|=4$ .

**Qwen3-4B-Thinking-2507:**

This is a complex or challenging question, and it is difficult to provide a direct and correct answer. I need to think about it. Well, so the problem is to find the largest possible real part of this expression:  $(75 + 117i)z + (96 + 144i)/z$ , where  $z$  is a complex number with  $|z| = 4$ . Hmm, let's start by recalling that when dealing with complex numbers on a circle of radius  $r$ , sometimes it's helpful to write  $z$  in polar form, like  $z = r e^{i\theta}$ , which here  $r = 4$ , so  $z = 4 e^{i\theta}$ . Then  $1/z$  would be  $(1/4) e^{-i\theta}$ , right? Maybe that substitution will make it easier to compute the real part. Let me try that first. Let's set  $z = 4(\cos\theta + i \sin\theta)$ , so  $1/z = (1/4)(\cos\theta - i \sin\theta)$ . Then let's compute each term separately and add them up, then take the real part. First term:  $(75 + 117i)z = (75 + 117i)(4 \cos\theta + 4i \sin\theta) = 4(75 + 117i)(\cos\theta + i \sin\theta)$ . Maybe I can factor out the 4 later, but let's just expand the product inside first. Let's compute  $(75 + 117i)(\cos\theta + i \sin\theta) = 75 \cos\theta + 75 i \sin\theta + 117 i \cos\theta + 117 i^2 \sin\theta = 75 \cos\theta - 117 \sin\theta + i(75 \sin\theta + 117 \cos\theta)$ . Okay, so then multiplying by 4 gives the first term as  $4(75 \cos\theta - 117 \sin\theta) + i4(75 \sin\theta + 117 \cos\theta)$ . Second term:  $(96 + 144i)/z = (96 + 144i)/(4 \cos\theta + 4i \sin\theta) = (24 + 36i)/(\cos\theta + i \sin\theta)$  [divided numerator and denominator by 4]. Maybe better to multiply numerator and denominator by the conjugate of the denominator to simplify? Wait, [...] **(omit 8459 tokens)** so we scale accordingly. Let's write:  $R(x, y) = 81x - 108y = 27(3x - 4y)$  Now define  $S = 3x - 4y$ . Then:  $\max_{x^2 + y^2 = 16} R(x, y) = 27 \cdot \max_{x^2 + y^2 = 16} S$  By the **Cauchy-Schwarz inequality**, for any real vectors  $\vec{u}, \vec{v}$ ,  $|\vec{u} \cdot \vec{v}| \leq |\vec{u}| \cdot |\vec{v}|$ . Here,  $\vec{u} = (3, -4)$ ,  $\vec{v} = (x, y)$ , and  $|\vec{v}| = 4$ . So:  $|S| = |3x - 4y| \leq \sqrt{3^2 + (-4)^2} \cdot 4 = 5 \cdot 4 = 20$ . Thus, the **maximum** value of  $S$  is 20, and hence:  $\max R(x, y) = 27 \cdot 20 = 540$  **Step 4: Confirm Attainability** Equality in Cauchy-Schwarz occurs when  $(x, y)$  is a scalar multiple of  $(3, -4)$ . Let:  $x = \frac{12}{5}$ ,  $y = -\frac{16}{5}$ . Check that  $x^2 + y^2 = (\frac{12}{5})^2 + (-\frac{16}{5})^2 = \frac{144 + 256}{25} = \frac{400}{25} = 16$ , which satisfies the constraint. Plug into  $R(x, y)$ :  $R = 81x - 108y = 81 \cdot \frac{12}{5} - 108 \cdot (-\frac{16}{5}) = \frac{972 + 1728}{5} = \frac{2700}{5} = 540$ . So the maximum is **attainable**. **Final Answer**  $\boxed{540}$   **(9497 tokens)**

**RPAM (ours):**


Okay, let's try to solve this problem step by step. The problem is to find the largest possible real part of the expression:  $(75 + 117i)z + \frac{96 + 144i}{z}$  where  $z$  is a complex number with  $|z| = 4$ . First, since  $|z| = 4$ , we can write  $z = 4e^{i\theta}$  for some real  $\theta$ . Alternatively, we can use the fact that if  $|z| = 4$ , then  $\frac{1}{z} = \frac{\overline{z}}{|z|^2} = \frac{\overline{z}}{16}$ , because  $z \overline{z} = |z|^2 = 16$ , so  $\frac{1}{z} = \frac{\overline{z}}{16}$ . That might be useful because it relates  $\frac{1}{z}$  to the conjugate of  $z$ , which could help in simplifying the expression. Let me denote  $z = x + yi$ , where  $x, y \in \mathbb{R}$ , and  $x^2 + y^2 = 16$  (since  $|z| = 4$ ). Then  $\frac{1}{z} = \frac{x - yi}{x^2 + y^2} = \frac{x - yi}{16}$ . So let's substitute into the expression: First term:  $(75 + 117i)(x + yi)$  Second term: [...] **(omit 5707 tokens)** Compute the real part directly: - First term:  $(75 + 117i)(x + iy) = 75x - 117y + i(75y + 117x)$  - Second term:  $\frac{96 + 144i}{16}(x + iy) = \frac{(96 + 144i)(x - iy)}{16}$  Real part of second term:  $\frac{1}{16}(96x + 144y) = 6x + 9y$  Total real part:  $(75x - 117y) + (6x + 9y) = 81x - 108y$  Maximize  $81x - 108y$  under  $x^2 + y^2 = 16$ . By Cauchy-Schwarz:  $(81x - 108y)^2 \leq (81^2 + 108^2)(x^2 + y^2) \Rightarrow |81x - 108y| \leq \sqrt{81^2 + 108^2} \cdot \sqrt{16}$  Compute: -  $81^2 = 6561$  -  $108^2 = 11664$  - Sum:  $6561 + 11664 = 18225$  -  $\sqrt{18225} = 135$  -  $\sqrt{16} = 4$  So maximum value is  $135 \cdot 4 = 540$  **Final Answer**  $\boxed{540}$   **(6484 tokens)**

Figure 6: For a challenge problem from AIME2024, RPAM is able to employ Long-CoT pattern to solve it, instead of directly generating the concise solution. Furthermore, RPAM costs only 6484 tokens to reach the correct answer rather than 9497 tokens in Qwen3-4B-Thinking.

**Question (GPQA - 87):**You tag the inner capsid proteins of rotavirus with a GFP and deliver them to mice via oral gavage. By in vivo imaging, you observe that some of the fluorescence was destroyed by the stomach acid, and some were entangled in the intestinal mucus before they reached the secondary lymphoid organs. You identify the fluorescent Peyer patches and isolate the proliferating cell population. You sequence their variable heavy chain gene and observe high variability. Which process is taking place?

- A. class switching recombination
- B. somatic hypermutation
- C. VDJ recombination
- D. complement activation

**Qwen3-4B-Thinking-2507:**

Okay, let's try to figure out this question. So, the experiment involves tagging the inner capsid proteins of rotavirus with GFP and delivering them to mice via oral gavage. Then they use in vivo imaging and notice that some fluorescence is destroyed by stomach acid, and some get tangled in intestinal mucus before reaching the secondary lymphoid organs. They find fluorescent Peyer patches and isolate the proliferating cells. When they sequence the variable heavy chain genes, they see high variability. [...] (omit 1940 tokens) This is a process that occurs in **activated B cells** within **germinal centers** of secondary lymphoid tissues (like Peyer patches) after antigen exposure. SHM introduces **point mutations** into the **variable regions** of the antibody genes, leading to **increased diversity** and **affinity maturation**. This directly explains the **high variability** observed in the variable heavy chain genes.- **C. VDJ recombination**: This is the **initial** process of B cell development in the **bone marrow**, where V, D, and J gene segments are recombined to generate the first diversity in the BCR. It occurs **before** antigen exposure and is **not** a process that happens in **proliferating cells in Peyer patches** after antigen delivery.- **D. Complement activation**: This is part of the **innate immune system** and involves the complement cascade. It is **not** related to B cell receptor gene rearrangement or variability in the variable heavy chain.### Why Somatic Hypermutation (B) is the Correct Answer:- The **proliferating cells** in the Peyer patches are **activated B cells** that have encountered the antigen (rotavirus proteins).- These cells are in the **germinal centers**, where **somatic hypermutation** is a well-established process.- SHM introduces **point mutations** into the variable heavy chain region, leading to **high variability** in the sequenced genes.- This is **not** the result of VDJ recombination (which occurs earlier in B cell development), nor of class switching or complement activation.### **Final Answer:** **B** ✓ (2379 tokens)

**RPAM (ours):**

Let's break this down step by step. We are given a scenario where: 1. **Rotavirus inner capsid proteins** are tagged with **GFP** and delivered to mice via **oral gavage**. 2. In vivo imaging shows: - Some fluorescence is destroyed by **stomach acid** (so some of the protein gets degraded in the stomach). - Some proteins get **entangled in intestinal mucus** before reaching secondary lymphoid organs. 3. The fluorescent **Peyer patches** (which are secondary lymphoid tissues in the gut) are identified, and the **proliferating cell population** is isolated. 4. When we sequence the **variable heavy chain gene** of these proliferating cells, we observe **high variability** [...] (omit 511 tokens) There is **no mechanism** for complement activation to cause **high variability in variable heavy chain genes**. → **Not D**### Final Analysis:- The key clue is: **proliferating cells in Peyer patches** show **high variability in variable heavy chain genes**. - This is classic **somatic hypermutation** — a process that occurs in activated B cells within germinal centers (such as in Peyer patches) to generate diversity in antibody affinity through point mutations in the variable region. Other processes either don't occur in this context (like VDJ recombination in mature B cells), or don't produce high variability in the variable region (like class switching).### ✓ Correct Answer: **B. somatic hypermutation** **B** ✓ **Final Answer:** **B** ✓ (814 tokens)

Figure 7: Even for an out-of-distribution multiple-choice problem from GPQA, RPAM is still able to generate Short-CoT response when the problem is relatively simple.