

AALC: Large Language Model Efficient Reasoning via Adaptive Accuracy-Length Control

Anonymous ACL submission

Abstract

Large reasoning models (LRMs) achieve impressive reasoning capabilities by generating lengthy chain-of-thoughts, but this “overthinking” incurs high latency and cost without commensurate accuracy gains. In this work, we introduce AALC, a lightweight, accuracy-aware length penalty integrated into reinforcement learning that dynamically balances correctness and brevity during training. By incorporating validation accuracy into the reward and employing a dynamical schedule mechanism, AALC delays the length penalty until target performance is met. Through extensive experiments across standard and out-of-distribution math benchmarks, we show that our approach reduces response length by over 50% while maintaining or even improving the original accuracy. Furthermore, qualitative analysis reveals that our method curbs redundant reasoning patterns such as excessive subgoal setting and verification, leading to structurally refined outputs rather than naive truncation. We also identify that efficiency gains are accompanied by reduced interpretability: models trained with AALC omit some narrative framing and explanatory context. These findings highlight the potential of reward-based strategies to guide LRMs toward more efficient, generalizable reasoning paths.¹

1 Introduction

Recent advances in Large reasoning models (LRMs) (OpenAI, 2024; Guo et al., 2025), have significantly improved performance on complex reasoning tasks by scaling computational effort at inference time. This enables models to produce exploratory reasoning chains resembling internal deliberation, facilitating self-assessment and correction (OpenAI, 2024; Guo et al., 2025; Team,

2024). Reinforcement learning (RL) has played a central role in enabling such capabilities. During RL training, models exhibit distinct reasoning phases characterized by longer, more intricate outputs and emergent strategies such as self-verification and decomposition (Gandhi et al., 2025). Empirical findings suggest a strong correlation between reasoning depth and accuracy gains (Zeng et al., 2025; Guo et al., 2025).

Despite their benefits, extended reasoning chains often introduce inefficiencies due to unnecessary repetition and elaboration—a phenomenon known as overthinking (Chen et al., 2024). This results in increased latency and computational cost with only marginal gains in accuracy. Motivated by Kimi1.5 (Team et al., 2025), which reports a reward function to shorten the chain-of-thought by using a length penalty, we proposed AALC reward to reduce models’ inference length while maintaining models’ performance.

This work presents two main contributions. We first introduce a lightweight reward mechanism that can be easily integrated into current reinforcement learning frameworks, featuring an accuracy-aware reward function that significantly reduces model reasoning length without compromising accuracy. Secondly, we provide an extensive evaluation across standard benchmarks to assess the trade-offs between response length and reasoning quality. Our method substantially reduces inference length while preserving both accuracy and semantic quality across standard and out-of-distribution math reasoning benchmarks, demonstrating the effectiveness and generality of our approach. In addition, our quantitative and qualitative analysis reveals:

1. For each benchmark, we observe a critical inflection point in average response length: accuracy improves or remains stable above this point, but declines sharply below it.
2. Behavior analysis reveals that excessive “sub-

^{*}Both authors contributed equally to this work.

¹Code is available at: <https://anonymous.4open.science/r/AALC-1B56>

goal setting” and “verification” primarily drive response verbosity; our method effectively suppresses the frequency of these and two additional reasoning behaviors.

3. When AALC is applied to both a base model and its distilled variant (e.g., Qwen2.5-Math and DeepSeek-R1-Distill-Qwen), the distilled model produces outputs with lengths comparable to those of the base model while maintaining accuracy.

2 Related Work

Efficient Reasoning Although increasing inference-time computational resources can improve the reasoning performance of LLMs, it often leads to overly verbose reasoning chains, resulting in higher computational costs and longer inference times (Sui et al., 2025). This “overthinking” phenomenon has motivated research into methods that encourage more concise and efficient reasoning. To address this, several strategies have been proposed. Model-based approaches, such as reinforcement learning with length-aware rewards, train models to balance correctness with brevity (detailed in the next paragraph). Supervised fine-tuning on variable-length Chain-of-Thought (CoT) data, as used in Cot-valve (Ma et al., 2025) and TokenSkip (Xia et al., 2025), similarly promotes compact reasoning. Prompt-based methods like Token-Budget (Han et al., 2024) and Chain of Drafts (Xu et al., 2025a) guide models to be concise by explicitly constraining response length. Output-oriented techniques, including Coconut (Hao et al., 2024), Softcot (Xu et al., 2025b), and Self-Calibration (Huang et al., 2025), reduce token overhead by encoding reasoning steps in latent form. Finally, dynamic reasoning frameworks (Liao et al., 2025; Ding et al., 2025) adapt reasoning depth based on real-time feedback to optimize efficiency.

RL with Length Control Design Recent work has applied reinforcement learning (typically PPO (Schulman et al., 2017)) with explicit length-based rewards to control the length of Chain-of-Thought (CoT) reasoning. Arora and Zanette (2025) proposed a reward scheme favoring shorter correct answers, while Yeo et al. (2025) introduced a Dirichlet-based Cosine Reward and length penalties to stabilize CoT growth. Kimi k1.5 (Team et al., 2025) incorporates a length penalty into its policy optimization to improve

long CoT activations and enable effective model merging. Some prior works either hurt models’ performance while reducing the inference length or do not reduce inference length as much as we do (Aggarwal and Welleck, 2025; Shen et al., 2025; Arora and Zanette, 2025). Our method differs by incorporating validation accuracy directly into the reward function via an accuracy-aware dynamic scheduling mechanism, enabling a balance between brevity and correctness. Additionally, since our method only alters the reward computation, it remains compatible with various training algorithms, including PPO (Schulman et al., 2017) and GRPO (Shao et al., 2024).

3 Methodology

Current language models’ outputs include much redundant content. However, a lightweight and effective length control mechanism is lacking to prevent overthinking. In this section, we introduce a lightweight, length-controlled reward designed to prevent overthinking without compromising model performance. Unlike most prior works that rely solely on the length of predictions and reference texts to regulate inference length, we propose a more principled reward design that incorporates validation accuracy to better balance performance and brevity. Section 3.1 introduces our length penalty formulation, which follows the standard structure adopted in prior works. Then, the accuracy is incorporated into the reward through a dynamic attention mechanism, effectively combining both objectives to prevent overthinking without sacrificing accuracy. Finally, Section 3.2 presents two dynamic scheduling strategies tailored for GRPO (Shao et al., 2024) training, allowing the reward signal to adapt for training.

3.1 Length Penalty

In the training process, accuracy remains the primary objective, particularly in the initial stage, which we refer to as the “warm-up” phase when the model is still learning to produce correct outputs. To maintain model performance, it is crucial to incorporate accuracy into the design of the reward function. During training, validation accuracy is the most reliable metric for evaluating model performance. As validation accuracy reaches its peak, the focus gradually shifts toward reducing generation length while maintaining ac-

curacy. Specifically, the length penalty is activated only when the validation accuracy meets a particular condition; otherwise, the model is encouraged to focus on improving accuracy before optimizing length. At the step i of the GRPO training process, the following equations define the length penalty and total reward for a prediction:

$$\begin{aligned} r_{\text{acc}} &= \frac{A_{\text{val}}}{A_{\text{target}}} \\ r_{\text{len}} &= \min\left(1, \frac{L_{\text{pred}}}{L_{\text{max}}}\right) \in [0, 1] \\ P_{\text{len}} &= \min(r_{\text{acc}}^\beta, r_{\text{len}}) \in [0, 1] \\ R &= R_{\text{acc}} - \alpha P_{\text{len}} \end{aligned}$$

where A_{val} and A_{target} denote the validation accuracy and the dynamically scheduled target accuracy at iteration i . L_{pred} is the length of the model’s predicted output, and L_{max} is the preset maximum allowed length; r_{acc} and r_{len} represent the normalized accuracy and length, respectively. P_{len} is the length penalty. A value close to 0 implies that the rear is inactive (no penalty), either because the output is sufficiently short or the model has not yet achieved enough accuracy to trigger length constraints. R_{acc} is the raw reward that only focuses on the correctness of a prediction but overlooks text length. It is defined as: $R_{\text{acc}} = \mathbb{I}(y_{\text{pred}}, y_{\text{gold}})$, where $R_{\text{acc}} = 1$ indicates that the prediction contains the correct answer and 0 otherwise. R is the total reward for a prediction, which is a weighted combination of the raw reward R_{acc} and the length penalty P_{len} .

The hyperparameter α controls the impact of the length reward. A large α will hurt models’ performance. It will be discussed in section 4.3. The hyperparameter β controls the sensitivity of the penalty to accuracy. A larger β delays the activation of the length penalty until higher accuracy is achieved, while a smaller β allows earlier enforcement of length constraints.

Instead of applying a hard threshold on accuracy, these formulations use a smooth transition to modulate the length penalty, ensuring continuous control that adapts dynamically as the model improves. Notably, the reward function serves two goals: (1) to guide the model toward accurate predictions when performance is suboptimal, and (2) to promote concise reasoning when accuracy is sufficiently high.

3.2 Dynamic Schedule

The parameter β in the length penalty determines when the penalty starts to take effect, but it relies on the target accuracy A_{target} . A well-calibrated A_{target} is essential to reduce overthinking while preserving model performance. A naive approach is to manually set A_{target} and apply the length penalty only after the model surpasses a fixed accuracy or training step threshold. However, this method is impractical, as neither the optimal accuracy nor the convergence point is known beforehand. Inspired by the momentum (Qian, 1999) method used in learning rate scheduling algorithms, we propose two dynamic length penalty scheduling strategies that adaptively set A_{target} during training.

Exponential Moving Average (EMA). This method updates the target accuracy by smoothing it toward the best validation accuracy seen so far:

$$\begin{aligned} A_{\text{val}}^{\text{max}} &= \max_{0 < i \leq t} A_{\text{val},i} \\ A_{\text{target},i} &= \max\left(\epsilon \cdot A_{\text{target},i-1} + (1 - \epsilon) \cdot A_{\text{val}}^{\text{max}}, A_{\text{val}}^{\text{max}}\right) \end{aligned}$$

where t is the current step index, and the ϵ controls the inertia of the target. The initial value of $A_{\text{target},i}$ is 1. When validation accuracy exceeds the previous target, the target is directly updated to prevent lag.

Potential Scheduling (PS). This strategy models the target accuracy as the current best validation accuracy plus a decaying potential:

$$\begin{aligned} A_{\text{val}}^{\text{max}} &= \max_{0 < i \leq t} A_{\text{val},i} \\ P_i &= \begin{cases} 1 - A_{\text{val},0} & i = 0 \\ \min(1 - A_{\text{val}}^{\text{max}}, \epsilon \cdot P_{i-1}) & i > 0 \end{cases} \\ A_{\text{target},i} &= A_{\text{val}}^{\text{max}} + P_i \end{aligned}$$

where t is the same as the above method, and ϵ controls how fast the potential decays, smaller values reduce P_i more aggressively over time. The min function ensures $A_{\text{target},i} \leq 1$ even when accuracy improves rapidly. This method ensures that A_{target} always stays above the best validation accuracy while gradually reducing the potential gap.

4 Experiments

4.1 Settings

Models and Datasets We conduct our training experiments on two commonly used and

Models	GSM8k			MATH		
	Accuracy \uparrow	Tokens \downarrow	CCA \uparrow	Accuracy \uparrow	Tokens \downarrow	CCA \uparrow
Qwen2.5-Math	93.21	439.04	14.8	62.27	577.94	10.1
Qwen2.5-Math-RS	97.17	193.73	23.5	73.01	395.31	17.6
Qwen2.5-Math-TS	88.40	160.54	20.2	58.48	290.16	14.4
Qwen2.5-Math-AALC	97.59	97.01	36.0	73.87	288.24	24.3
DS-Distill-Qwen	94.72	455.61	21.0	70.33	1006.94	13.2
DS-Distill-Qwen-RS	97.06	497.88	33.7	82.87	557.12	25.8
DS-Distill-Qwen-AALC	97.72	100.58	52.2	82.77	270.51	38.6

Table 1: Performance of models on GSM8k and MATH datasets. Each model is trained and evaluated on the same dataset, and then assessed by three metrics. “RS” denotes that the base model is fine-tuned using raw score. “TS” refers to the model derived from the TokenSkip. “AALC” indicates that AALC reward is applied to the base model during the training process. “Qwen2.5-Math” and “DS-Distill-Qwen” are the abbreviations of “Qwen2.5-Math-7B” and “DeepSeek-R1-Distill-Qwen-7B”.

Models	AIME24		AMC24		CNMO24		GPQA	
	Accuracy \uparrow	Tokens \downarrow	Accuracy \uparrow	Tokens \downarrow	Accuracy \uparrow	Tokens \downarrow	Accuracy \uparrow	Tokens \downarrow
Qwen2.5-Math	6.7	1868.50	4.3	1509.89	11.1	1072.56	30.1	847.97
Qwen2.5-Math-RS	26.7	751.63	21.7	729.52	22.2	847.22	32.3	639.25
Qwen2.5-Math-AALC	26.7	421.51	23.9	395.10	22.2	546.67	30.0	568.01
DS-Distill-Qwen	16.7	3218.06	26.1	3142.20	16.7	3108.72	13.1	3054.42
DS-Distill-Qwen-RS	26.7	3038.67	47.8	1598.24	61.1	2887.39	32.3	3002.72
DS-Distill-Qwen-AALC	30.0	2724.07	52.2	1264.65	66.6	2729.83	37.9	2749.09

Table 2: Evaluate models on out-of-distribution math datasets.

well-known models: Qwen2.5-Math-7B (Yang et al., 2024) and DeepSeek-R1-Distill-Qwen-7B (DeepSeek-AI, 2025), which have shown excellent reasoning ability on various math tasks. Both models are trained on the GSM8k (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021) datasets. Both training set has about 7.5k questions. During our training, we split both training sets in 95%:5% ratio to create validation sets for the training process. We evaluate our models on different reasoning datasets: GSM8k, MATH, AIME24, AMC24, CNMO24, and GPQA (Rein et al., 2024). We fine-tune models using the GRPO methods. The default parameters during the training process are as follows. The values of α , β , and ϵ are 10^{-6} , 128, and 0.9, respectively. We choose the first dynamic scheduling method as the default training method. The length limitation of both prompts and generations is 1k tokens, which is sufficient for most of our questions and answers in training datasets. Models are fine-tuned for about 300 steps to reach convergence on both inference length and accuracy. All experiments are conducted on the AWS EC2 platform using 8 NVIDIA H100 GPUs.

Baselines To assess the effectiveness of our approach, we compare it against three baseline methods: (1) Original models, which are the publicly released models used without any additional fine-tuning. (2) Original models-RS, which are the original models fine-tuned using raw binary supervision. Specifically, if the answer extracted from the model’s generated reasoning is correct, a score of 1 is assigned; otherwise, a score of 0 is given. (3) Prior method: which refers to the TokenSkip (Xia et al., 2025) approach implemented using the same base models as in our experiments.

Metrics We evaluate using accuracy, average inference length, and Consistent Concise Accuracy (CCA) (Nayab et al., 2024), a composite metric that considers both efficiency and accuracy. For CCA, we follow their settings of $\alpha = 10$ and $\beta = 40$, and use $k = 1024$ in our experiments. All metrics are averaged over 3 runs.

4.2 Results

Table 1 demonstrates the performance of various methods on different models. Models trained with our rewards answer questions in short lengths,

while slightly improving accuracy. The TokenSkip model is evaluated only on the Qwen2.5-Math series, as no other trained versions are available.

Performance Analysis Notably, the results demonstrate that our length penalty framework can significantly reduce overthinking while improving the models’ reasoning ability compared to the original model. On the GSM8k and MATH datasets, models finetuned with AALC only need less than a quarter and a half of the original length to generate correct answers, respectively, and show a significant improvement in CCA. Moreover, applying the raw score (0/1 score) during training also causes the model outputs to gradually decrease in most cases. This is mainly because post-training can slightly simplify outputs by removing redundant content. Detailed analysis of the change of inference texts will be discussed in section 5. Furthermore, the training process on the GSM8k dataset is unstable, as shown in Figure 1. This is mainly because the dataset is simple, and post-training with raw scores as the reward cannot effectively simplify it.

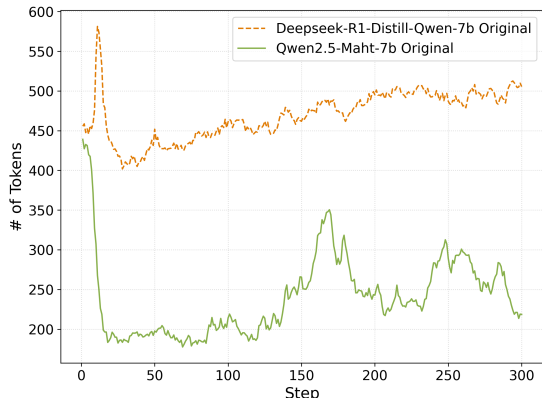


Figure 1: # of inference tokens of models during training with raw scores.

Evaluation on Out-of-Distribution Datasets

In addition to the above two datasets, we also evaluate our models trained on the MATH dataset, on out-of-distribution datasets, as shown in Table 2. The original models perform poorly due to constraints on generation length. Most outputs are truncated by the length limitation rule and do not contain correct answers. In contrast, models trained with AALC consistently outperform both the original models and those trained with raw scores across all datasets, while also generating fewer tokens. Thus, it demonstrates that our length penalty improves models’ generalizability and robustness on out-of-distribution complex datasets.

Scaling Analysis As shown in Figure 2 and Table 3, our AALC reward method consistently improves training across various model scales. All training runs incorporating AALC effectively mitigate overthinking while preserving accuracy compared to those using only raw scores (RS). Since the smallest official model in the Qwen2.5-Math family is 1.5B, we include Qwen2.5-0.5B as a substitute to examine behavior at a smaller scale, Although its training is less stable due to limited reasoning capability, it still follows the same performance trend observed in larger models, validating the generality of our approach.

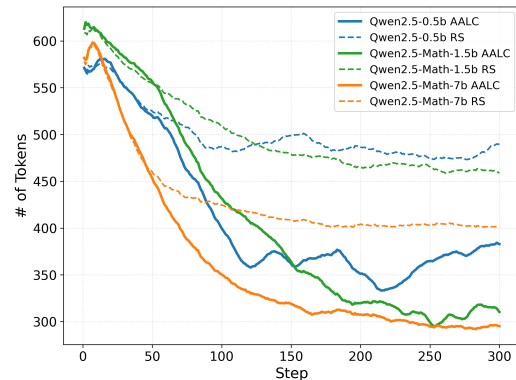


Figure 2: Training Progress of Qwen2.5 Models Across Scales. Bold lines indicate training runs that incorporate the AALC. Models of the same scale are represented with similar colors.

Models	Acc (RS)	Acc (AALC)
Qwen2.5-0.5b	34.0	33.9
Qwen2.5-Math-1.5b	67.8	68.4
Qwen2.5-Math-7b	73.0	73.9

Table 3: Accuracy comparison of Qwen2.5 models trained by Raw Score (RS) vs. AALC across scales.

4.3 Ablation Study

Hyper Parameters α , β , and ϵ We run all the following experiments following the second dynamic schedule method on Qwen2.5-Math-7B model training on the MATH dataset. The baseline settings are the same as the “Length Penalty” setting in the main result section, which is $\alpha = 1e^{-6}$, $\beta = 128$, and $\epsilon = 0.9$.

We evaluate five values of α : 10^{-5} , 5×10^{-6} , 2×10^{-6} , 10^{-6} , and 10^{-7} (see Figures 3b and 3e). Among them, $\alpha = 10^{-6}$ yields the best trade-off, effectively reducing overthinking while preserving high accuracy. A smaller value like 10^{-7} fails to curb verbosity, while a larger value such as 10^{-5} leads to a notable drop in accuracy (from 73% to

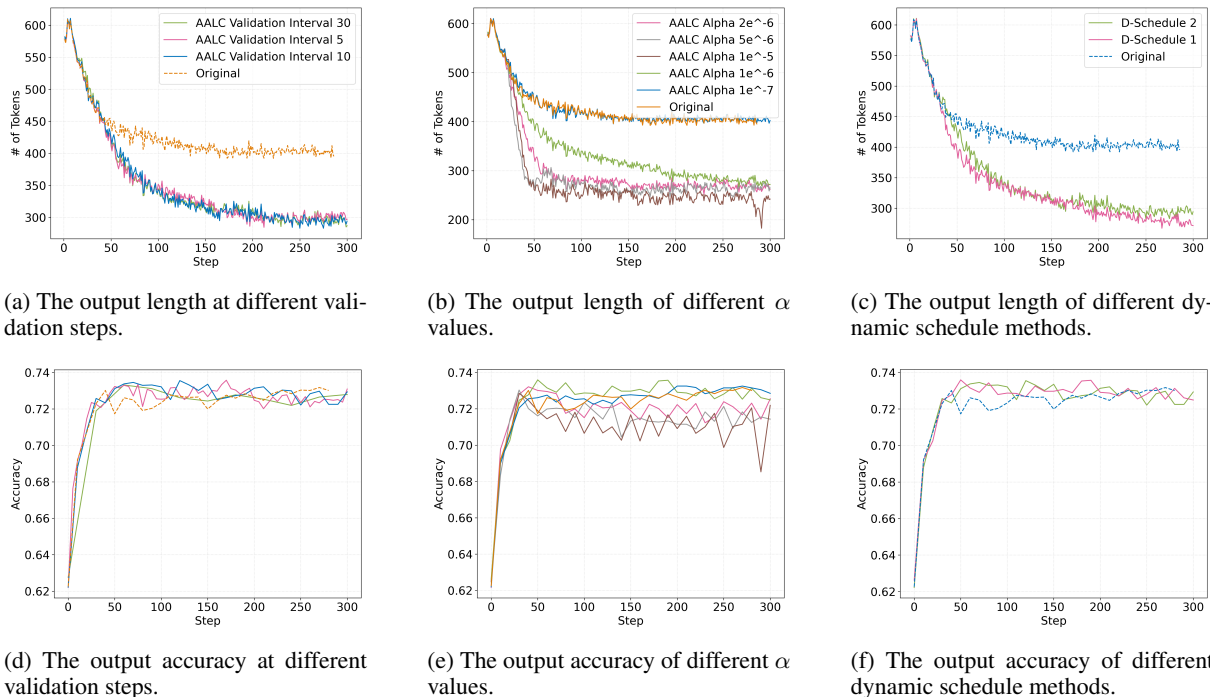


Figure 3: Ablation analysis of token usage and accuracy under different training configurations. Top row (a–c): Impact on output length (token count) when varying validation step interval (a), length reward weight α (b), and different dynamic schedule strategy (c). Bottom row (d–f): Corresponding effects on validation accuracy for the same settings. D-Schedule-1 and 2 refer to EMA and PS, respectively.

382 71%) with only a modest reduction of ~ 30 tokens
 383 in output length. Notably, the training curve for
 384 $\alpha = 10^{-5}$ shows greater fluctuation than others;
 385 both accuracy and length decline sharply toward
 386 the end of training before recovering. This pattern
 387 highlights the auto-adaptability of our dynamic
 388 schedule: when accuracy degrades, the method
 389 prioritizes restoring performance over enforcing
 390 brevity.

391 **Dynamic Schedule** We compare our dynamic
 392 scheduling methods under the same hyperparam-
 393 eter settings as shown in the above section. Figure
 394 3c and 3f show that the two methods are similar in
 395 effectiveness. However, the first dynamic schedul-
 396 ing method converges faster and more stably at the
 397 end of training.

398 **Validation Interval** Since validation accuracy is
 399 an essential part of formulas and controls two parts
 400 of the framework, the frequency of obtaining the
 401 value may be vital to the training process. We have
 402 tried three interval values: 5, 10, and 30. Results
 403 are in Figure 3a and 3d. Experiments show that
 404 there is almost no difference in choosing different
 405 validation intervals. With a small validation inter-
 406 val, the reduction of generation length is smoother
 407 than with others. However, it does not affect the

408 final results. More detailed ablation study results
 409 are in Appendix C.

5 Analysis 410

5.1 Semantic Quality Analysis 411

412 In this section, in addition to standard metrics
 413 such as accuracy and response length, we evalu-
 414 ate the semantic quality of reduced responses,
 415 providing a complementary dimension for assess-
 416 ing efficiency in reasoning tasks. Specifically, we
 417 present the win rate of our proposed method com-
 418 pared to other baselines in Table 1. This evaluation
 419 includes both manual assessments and automated
 420 comparisons using GPT-4o (Hurst et al., 2024).
 421 During these pairwise comparisons, each model
 422 response earns 1 point for a win, 0.5 points for a
 423 tie, and 0 points for a loss. The win rate is thus cal-
 424 culated as the proportion of total points earned by
 425 our method relative to the baseline method across
 426 all comparisons. Given a problem instance x with
 427 corresponding solutions y_1 and y_2 , the evaluation
 428 criteria are as follows:

- 429 • If y_1 is correct and y_2 is not, y_1 wins; if y_2 is
 430 correct and y_1 is not, y_2 wins;
- 431 • If both y_1 and y_2 are incorrect, neither re-
 432 ceives points;

Our methods	Opponent	GSM8k		MATH	
		Human	GPT-4o	Human	GPT-4o
Qwen2.5-Math-AALC	Qwen2.5-Math	64.4%	55.9%	66.8%	57.2%
	Qwen2.5-Math-RS	60.5%	53.4%	58.3%	53.6%
DS-Distill-Qwen-AALC	DS-Distill-Qwen	61.3%	52.8%	67.5%	54.6%
	DS-Distill-Qwen-RS	57.8%	52.5%	61.6%	54.1%

Table 4: Win rates of our proposed method against various baselines settings on GSM8k and MATH benchmark. Evaluations include human annotators and GPT-4o.

- If both are correct, their semantic quality is evaluated by human or GPT-4o to determine the superior one.

The evaluation prompt template, detailed in Appendix B, is specifically designed to mitigate length bias during semantic quality assessment. To avoid positional bias, the order of y_1 and y_2 is randomized. This evaluation approach is motivated by the fact that, in reasoning tasks, only correct solutions are semantically meaningful; incorrect solutions lack value regardless of their conciseness or perceived faithfulness.

The results of this analysis are presented in Table 4. Our approach consistently achieves a win rate above 50% against all baselines. Moreover, we observe that human evaluations yield a higher win rate compared to GPT-4o assessments, likely due to GPT-4o’s known preference bias towards lengthier responses (Singhal et al., 2023). Despite a slight reduction in raw accuracy scores compared to baseline methods shown in Table 1, our method demonstrates superior semantic quality. These findings validate our method’s ability to generate concise yet semantically richer solutions.

5.2 Reasoning Behavior Analysis

To understand the changes of model’s reasoning patterns throughout the training process, following (Zeng et al., 2025), we adopt the cognitive behavior framework proposed by Gandhi et al. (2025), leveraging GPT-4o (Hurst et al., 2024) to identify four behaviors: “Backtracking”, “Verification”, “Subgoal Setting”, and “Enumeration”. This approach is more robust compared to traditional keyword-based methods (Xie et al., 2025; Yeo et al., 2025), effectively capturing nuanced reasoning patterns even without explicit textual indicators. Definitions and evaluation prompts for these behaviors are provided in Appendix A.

Specifically, following the experiment setting in Table 1, we conduct this reasoning behavior anal-

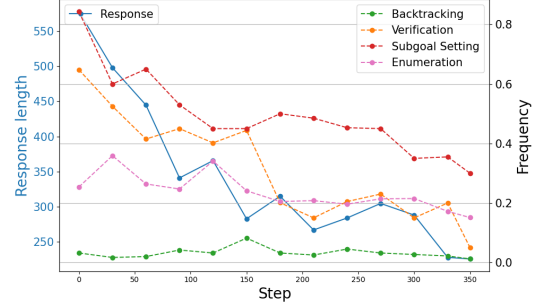


Figure 4: Response length and frequency of reasoning behaviors identified in Qwen2.5-Math-AALC on MATH.

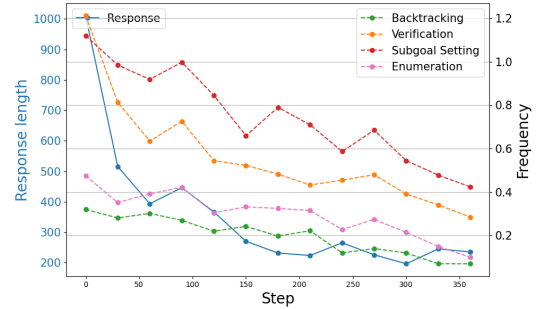


Figure 5: Response length and frequency of reasoning behaviors identified in DS-Distill-Qwen-AALC on MATH.

ysis on two models, Qwen2.5-Math-AALC and DS-Distill-Qwen-AALC on the MATH benchmark. Figures 4 and 5 show the correlation between response length and the frequency of identified reasoning behaviors during training. Our analysis reveals that “Subgoal Setting” and “Verification” are predominant, significantly contributing to reasoning verbosity. As training progresses, we observe a concurrent reduction in the frequency of all identified reasoning behaviors, aligning closely with decreased output lengths. These observations suggest that our method for improving reasoning efficiency is primarily achieved by reducing redundant reasoning behavior.

Dataset / Question	Output of Raw Score	Output of AALC
MATH: How many vertical asymptotes does the graph of $y = \frac{2}{x^2+x-6}$ have?	To determine the number of vertical asymptotes of the function $y = \frac{2}{x^2+x-6}$, we need to find the values of x where the denominator is zero. The denominator is $x^2 + x - 6 = (x + 3)(x - 2)$. So, the solutions are $x = -3$ and $x = 2$. These are the vertical asymptotes. Final answer: 2 .	The function $y = \frac{2}{x^2+x-6}$ has vertical asymptotes where the denominator is zero. The denominator x^2+x-6 factors as $(x+3)(x-2)$. So, the asymptotes are at $x = -3$ and $x = 2$. Final answer: 2 .
GSM8k: A robe takes 2 bolts of blue fiber and half that much white fiber. How many bolts in total does it take?	The robe takes 2 bolts of blue fiber and half that, i.e., 1 bolt, of white fiber. Total = $2 + 1 = 3$.	2 blue bolts + 1 white bolt = 3 bolts in total. Final answer: 3 .

Table 5: Compare the outputs of Qwen2.5-Math model trained with raw score and that trained with dynamic length penalty ($\alpha = 1e-6$).

5.3 More Qualitative Analysis

To understand the impact of our AALC reward on inference behavior, we compare the Qwen2.5-Math models trained with Raw Score with a variant trained with AALC (with a learning rate of $1e-6$). A case study is provided in Table 5. While both models yield correct answers, we observe that the length-penalized model significantly compresses explanations and eliminates scaffolding, presenting solutions in a more direct and streamlined manner. This behavior is consistent across both the MATH and GSM8k datasets. Structural elements such as framing phrases, code-style reasoning, and boxed outputs are removed entirely. These shorter, higher-confidence completions often result in improved exact match accuracy. We attribute these gains to the model’s preference for more direct solution paths under length constraints, which suppresses narrative redundancy and promotes structurally compressed, low-variance reasoning.

More Insights Further analysis of early- vs. mid-training stages and base vs. distilled models reveals three key insights (Appendix D): First, the effects of length penalty emerge very early (by training step 60, see Appendix D.1). Second, gains in efficiency coincide with a loss of interpretability: model outputs increasingly omit narrative framing, explanatory context, and derivation steps (see Table 6). Third, when AALC is applied to both a base model and its distilled version (e.g.,

Qwen2.5-Math and DeepSeek-R1-Distill-Qwen), the distilled model’s outputs converge to lengths comparable to the base model’s while preserving accuracy (see Appendix D.2). These findings suggest that the benefit of reduced output length arise from structural refinements in reasoning process rather than mere truncation.

6 Conclusion

In this work, we present a lightweight, accuracy-aware length penalty mechanism to mitigate overthinking in large language models, promoting concise yet accurate reasoning. By dynamically balancing brevity and correctness within a reinforcement learning framework, our method consistently reduces response length, often by over 50%, while preserving or even improving performance. These results hold consistently across model families, scales, and out-of-distribution datasets, highlighting the generality of our approach. Analyses reveal that our approach suppresses redundant cognitive behaviors, such as excessive sub-goal setting and verification, and induces structural refinement rather than mere truncation. These findings demonstrate the potential of reward-based strategies to guide LLMs toward more efficient and robust reasoning paths.

Limitation Our approach is mainly tested on math benchmarks and its generalization to tasks needing richer language, such as common-sense reasoning or dialog, is unverified. The

487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517

518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547

548	method’s effectiveness is sensitive to hyperparameter choices, and does not dynamically adjust reasoning depth based on task complexity, which may limit adaptability. Additionally, evaluation partly depends on GPT-4o, which may introduce bias despite mitigation efforts.	
549		
550		
551		
552		
553		
554	References	
555	Pranjal Aggarwal and Sean Welleck. 2025. L1: Controlling how long a reasoning model thinks with reinforcement learning. <i>arXiv preprint arXiv:2503.04697</i> .	
556		
557		
558		
559	Daman Arora and Andrea Zanette. 2025. Training language models to reason efficiently. <i>arXiv preprint arXiv:2502.04463</i> .	
560		
561		
562	Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, et al. 2024. Do not think that much for $2+3=?$ on the overthinking of o1-like llms. <i>arXiv preprint arXiv:2412.21187</i> .	
563		
564		
565		
566		
567	Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. <i>arXiv preprint arXiv:2110.14168</i> .	
568		
569		
570		
571		
572	DeepSeek-AI. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning .	
573		
574	Yifu Ding, Wentao Jiang, Shunyu Liu, Yongcheng Jing, Jinyang Guo, Yingjie Wang, Jing Zhang, Zengmao Wang, Ziwei Liu, Bo Du, et al. 2025. Dynamic parallel tree search for efficient llm reasoning. <i>arXiv preprint arXiv:2502.16235</i> .	
575		
576		
577		
578		
579	Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. 2025. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. <i>arXiv preprint arXiv:2503.01307</i> .	
580		
581		
582		
583		
584	Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and Noah D Goodman. 2024. Stream of search (sos): Learning to search in language. <i>arXiv preprint arXiv:2404.03683</i> .	
585		
586		
587		
588		
589	Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. <i>arXiv preprint arXiv:2501.12948</i> .	
590		
591		
592		
593		
594		
595	Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. 2024. Token-budget-aware llm reasoning. <i>arXiv preprint arXiv:2412.18547</i> .	
596		
597		
598		
	Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. Training large language models to reason in a continuous latent space. <i>arXiv preprint arXiv:2412.06769</i> .	599 600 601 602 603
	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. <i>arXiv preprint arXiv:2103.03874</i> .	604 605 606 607 608
	Chengsong Huang, Langlin Huang, Jixuan Leng, Jiacheng Liu, and Jiaxin Huang. 2025. Efficient test-time scaling via self-calibration. <i>arXiv preprint arXiv:2503.00031</i> .	609 610 611 612
	Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. <i>arXiv preprint arXiv:2410.21276</i> .	613 614 615 616 617
	Baohao Liao, Yuhui Xu, Hanze Dong, Junnan Li, Christof Monz, Silvio Savarese, Doyen Sahoo, and Caiming Xiong. 2025. Reward-guided speculative decoding for efficient llm reasoning. <i>arXiv preprint arXiv:2501.19324</i> .	618 619 620 621 622
	Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. 2025. Cot-valve: Length-compressible chain-of-thought tuning. <i>arXiv preprint arXiv:2502.09601</i> .	623 624 625 626
	Sania Nayab, Giulio Rossolini, Marco Simoni, Andrea Saracino, Giorgio Buttazzo, Nicolamaria Manes, and Fabrizio Giacomelli. 2024. Concise thoughts: Impact of output length on llm reasoning and cost. <i>arXiv preprint arXiv:2407.19825</i> .	627 628 629 630 631
	OpenAI. 2024. Learning to reason with llms . Accessed: 2025-04-30.	632 633
	Ning Qian. 1999. On the momentum term in gradient descent learning algorithms . <i>Neural Networks</i> , 12(1):145–151.	634 635 636
	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 <i>First Conference on Language Modeling</i> .	637 638 639 640 641
	John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. <i>arXiv preprint arXiv:1707.06347</i> .	642 643 644 645
	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. 2024. Deepseek-math: Pushing the limits of mathematical reasoning in open language models. <i>arXiv preprint arXiv:2402.03300</i> .	646 647 648 649 650 651

652 Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi,
653 Wenjing Zhang, Jiangze Yan, Ning Wang, Kai
654 Wang, and Shiguo Lian. 2025. Dast: Difficulty-
655 adaptive slow-thinking for large reasoning models.
656 *arXiv preprint arXiv:2503.04472*.

657 Prasann Singhal, Tanya Goyal, Jiacheng Xu, and Greg
658 Durrett. 2023. A long way to go: Investigat-
659 ing length correlations in rlhf. *arXiv preprint*
660 *arXiv:2310.03716*.

661 Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu
662 Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, An-
663 drew Wen, Hanjie Chen, Xia Hu, et al. 2025.
664 Stop overthinking: A survey on efficient reason-
665 ing for large language models. *arXiv preprint*
666 *arXiv:2503.16419*.

667 Kimi Team, Angang Du, Bofei Gao, Bowei Xing,
668 Changju Jiang, Cheng Chen, Cheng Li, Chenjun
669 Xiao, Chenzhuang Du, Chonghua Liao, et al. 2025.
670 Kimi k1. 5: Scaling reinforcement learning with
671 llms. *arXiv preprint arXiv:2501.12599*.

672 Qwen Team. 2024. [Qwq-32b-preview](#). Accessed:
673 2025-04-30.

674 Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi
675 Li, and Wenjie Li. 2025. [TokenSkip: Control-
676 lable chain-of-thought compression in LLMs](#). pages
677 3351–3363.

678 Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo,
679 Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhi-
680 rong Wu, and Chong Luo. 2025. Logic-rl: Un-
681 leashing llm reasoning with rule-based reinforc-
682 ement learning. *arXiv preprint arXiv:2502.14768*.

683 Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng
684 He. 2025a. Chain of draft: Thinking faster by writ-
685 ing less. *arXiv preprint arXiv:2502.18600*.

686 Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan
687 Miao. 2025b. Softcot: Soft chain-of-thought
688 for efficient reasoning with llms. *arXiv preprint*
689 *arXiv:2502.12134*.

690 An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao,
691 Bowen Yu, Chengpeng Li, Dayiheng Liu, Jian-
692 hong Tu, Jingren Zhou, Junyang Lin, Keming Lu,
693 Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang
694 Ren, and Zhenru Zhang. 2024. Qwen2.5-
695 math technical report: Toward mathematical ex-
696 pert model via self-improvement. *arXiv preprint*
697 *arXiv:2409.12122*.

698 Edward Yeo, Yuxuan Tong, Morry Niu, Graham Neu-
699 big, and Xiang Yue. 2025. Demystifying long
700 chain-of-thought reasoning in llms. *arXiv preprint*
701 *arXiv:2502.03373*.

702 Weihao Zeng, Yuzhen Huang, Qian Liu, Wei Liu,
703 Keqing He, Zejun Ma, and Junxian He. 2025.
704 Simplerl-zoo: Investigating and taming zero rein-
705 forcement learning for open base models in the wild.
706 *arXiv preprint arXiv:2503.18892*.

A Reasoning Behavior Analysis

707

Gandhi et al. (2024) identify four core cognitive behaviors displayed by language models:

708

- (1) **Verification:** systematically checking intermediate results for correctness;
- (2) **Backtracking:** detecting errors mid-generation and explicitly revising prior steps;
- (3) **Subgoal setting:** decomposing a complex problem into smaller, tractable sub-tasks;
- (4) **Enumeration:** exhaustively considering multiple cases or possibilities when searching for a solution.

709

710

711

712

713

We substitute “Enumeration” for the original “Backward Chaining,” which is less pertinent to mathematical-reasoning tasks.

714

715

Prompt Template for Identifying and Analyzing Reasoning Behaviors

Below is a chain-of-reasoning generated by a Language Model when attempting to solve a math problem. Evaluate this chain-of-reasoning to determine whether it demonstrates beneficial problem-solving behaviors that deviate from typical linear, monotonic reasoning patterns commonly observed in language models.

```
<start_of_reasoning>
{input}
<end_of_reasoning>
```

Specifically, actively identify and emphasize beneficial behaviors such as:

- **Backtracking:** Explicitly revising approaches upon identifying errors or dead ends (e.g., *"This approach won't work because..."*).
- **Verification:** Systematically checking intermediate results or reasoning steps (e.g., *"Let's verify this result by..."*).
- **Subgoal Setting:** Breaking down complex problems into smaller, manageable steps (e.g., *"To solve this, we first need to..."*).
- **Enumeration:** Solving problems by exhaustively considering multiple cases or possibilities.

Additionally, remain attentive to and encourage the identification of other beneficial behaviors not explicitly listed here, such as creative analogies, abstraction to simpler cases, or insightful generalizations.

Important:

Clearly specify each beneficial behavior you identify.

Provide explicit examples from the reasoning chain.

If no beneficial behaviors are observed, explicitly return an empty list.

Provide your evaluation clearly, formatted as follows:

```
{
  "behavior": "",
  "example": ""
}
```

716

A.1 Qwen3-235B-A22B Output

Qwen Output

We are given:

- Gloria can either buy a **pair of boots** or **two pairs of high heels**. - The **total cost of the two pairs of heels is \$5 less than the cost of the boots**. - One pair of heels costs \$33, and the other costs twice as much.

Step 1: Total cost of heels:

$$33 + (2 \times 33) = 33 + 66 = 99$$

Step 2: Cost of boots:

$$99 + 5 = \boxed{104}$$

A.2 ChatGPT-4o Output

Qwen Output

We are given:

- Gloria can either buy a **pair of boots** or **two pairs of high heels**.
- The **total cost of the two pairs of heels is \$5 less than the cost of the boots**.
- One pair of heels costs \$33, and the other costs **twice as much**.

Step 1: Define variables

Let:

- B = cost of the **boots**
- The first pair of heels = \$33
- The second pair of heels = $\$2 \times 33 = \66

So, total cost of both pairs of heels =

$$33 + 66 = 99$$

We're told this is \$5 less than the boots, so:

$$B = 99 + 5 = 104$$

Final Answer:

The boots cost **\$104**.

ChatGPT Output

Let's solve the problem **step by step**.

Step 1: Understand the problem

Gloria finds a pair of **boots** that cost a certain amount. She also considers buying **two pairs of high heels** instead.

- The **first pair** of heels costs \$33.
- The **second pair** of heels costs **twice as much**:

$$2 \times 33 = 66 \text{ dollars}$$

- The **total cost of the two heels** is:

$$33 + 66 = 99 \text{ dollars}$$

- The **boots cost \$5 more** than the two pairs of heels:

$$\text{Boots} = 99 + 5 = 104 \text{ dollars}$$

Final Answer:

104

Gemini-2.0 Output

Gemini Output

Let's break this down step by step:

- 1. Cost of the more expensive heels:** One pair of heels costs \$33, and the other costs twice as much. So, the more expensive heels cost $\$33 \times 2 = 66$.
- 2. Total cost of the two pairs of heels:** To find the combined cost of both pairs of heels, we add their individual prices: $\$33 + 66 = 99$.
- 3. Cost of the boots:** The problem states that the two pairs of high heels together cost five dollars less than the boots. This means the boots cost \$5 more than the combined cost of the heels. So, the boots cost $\$99 + 5 = 104$.

Therefore, the boots cost **104** dollars.

DeepSeek Output

DeepSeek Output

Let's determine the cost of the boots step by step.

1. Identify the Costs of the High Heels:

- **First pair of heels:** \$33
- **Second pair of heels:** Twice as much as the first pair.

$$2 \times \$33 = \$66$$

2. Calculate the Total Cost of Both Pairs of Heels:

$$\$33 \text{ (first pair)} + \$66 \text{ (second pair)} = \$99$$

3. Determine the Cost of the Boots:

- The problem states that the two pairs of heels together cost **\$5 less** than the boots.
- Let B be the cost of the boots.

$$\$99 = B - \$5$$

$$B = \$99 + \$5 = \$104$$

Final Answer:

104

B Win Rate Analysis

727

We use the following prompt template for win rate analysis with three criteria that helps GPT-4o keep a focus on instruction following and semantic quality while avoiding length bias:

728

729

Response Semantic Quality Comparison Template

Given the following instruction and two candidate responses:

Instruction: {}

Response 1: {}

Response 2: {}

Please evaluate which response better follows the instruction based on the following criteria:

1. Adherence to the instruction.
2. Overall semantic quality.
3. Conciseness: the response should **avoid** redundancy, unnecessary verbosity, or inclusion of irrelevant information.

Based on your evaluation, respond in the format:

'Response 1 is better than Response 2', 'Response 2 is better than Response 1', or 'Response 1 is equal to Response 2'.

730

C More Ablation Study

731

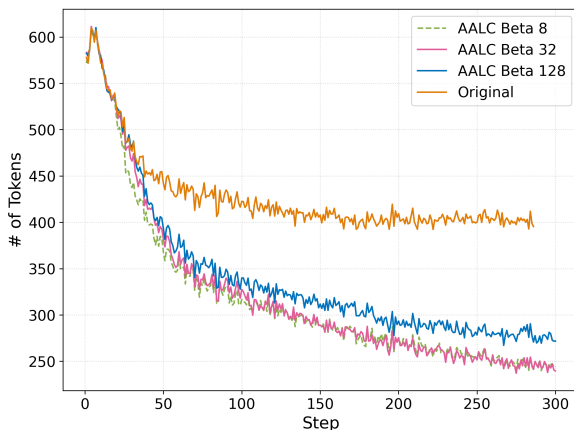


Figure 6: The output length at different β values.

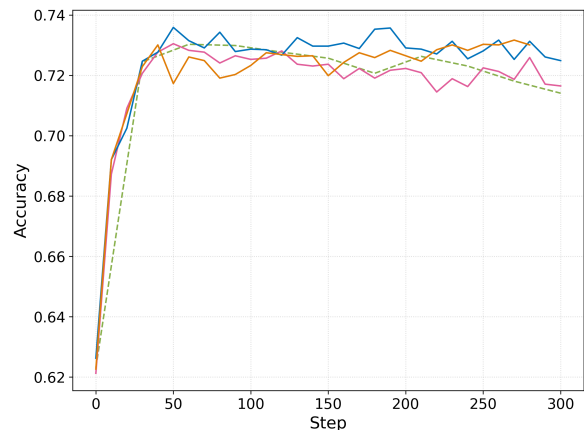


Figure 7: The output accuracy at different β values.

For β , we have tested the following values: 8, 32, and 128. Results are in the figures 6 and 7. The β controls when the length penalty intervenes in the training process. With a low β value, such as 8, the length penalty becomes effective early, as the validation accuracy reaches around 45% of the target accuracy. According to Figure 6, early engagement in the training process can reduce the generation length more, but the model's performance also decreases from 73% to 71%. With a high β value, such as 128, the length penalty does not affect the training process until the validation accuracy reaches 95% of the target accuracy. The validation accuracy process also shows the model's performance is even better than the "Original" training process. In our framework, our goal is to maintain the model's performance while preventing overthinking. Thus, a high value is preferred in this case.

732

733

734

735

736

737

738

739

740

D More Qualitative Analysis

Dataset / Question	Step 30 Checkpoint Output	Step 60 Checkpoint Output
Math: <i>What is the sum of the first 10 positive even numbers?</i>	Detailed explanation: defines sequence, uses formula $S = \frac{n}{2}(a + l)$, explains each term, then computes: $S = 10/2 * (2 + 20) = 5 * 22 = \boxed{110}$ Includes full derivation and justification (not shown here).	Condensed version: <i>“First 10 even numbers: 2 to 20. Sum = $10/2 * (2 + 20) = 5 * 22 = 110$. ##### 110”</i> No sequence explanation or setup.
GSM8K: <i>Kylar wants to buy 16 glasses. One costs \$5; every second glass costs 60%.</i>	Full arithmetic + Python code: <pre>cost_first = 5, cost_second = 5 * 0.6 total_cost = (cost_first + cost_second) * 8 print(total_cost) # Output: 64.0</pre> Explains pairing logic and includes validation (omitted here).	Shortened to only arithmetic steps: <i>“8 cheaper glasses cost $8 * \\$3 = \\24. 8 full-price glasses cost $8 * \\$5 = \\40. Total = 64. ##### 64”</i> All narrative and code removed.

Table 6: Comparison between early (Step 30) and mid-stage (Step 60) model outputs. Length penalty begins to reduce verbosity while preserving core reasoning accuracy.

742

D.1 When Does the AALC Start Working? A Step 30 vs Step 60 Analysis

743

744

745

746

747

748

749

To investigate when the AALC begins to meaningfully affect model behavior, we compare inference outputs at step 30 and step 60 under AALC settings (1e-6 learning rate, dynamic schedule) on both the MATH and GSM8k datasets. These checkpoints are chosen to align with the transition period of the d-schedule mechanism, which gradually increases the weight of the length penalty during post-training. By step 60, we already observe clear signs of compression: outputs are more concise, pedagogical scaffolding is reduced, and verbose explanation structures are stripped away. Table 6 illustrates this behavioral shift with representative examples.

750

751

752

753

754

755

756

757

At step 30, the model still retains a “teacher-like” reasoning style. For example, in response to a math problem asking for the sum of the first 10 positive even numbers, the model describes the arithmetic sequence, defines variables explicitly, applies the formula $S = \frac{n}{2}(a + l)$, and concludes with a boxed answer. In contrast, the step 60 output removes the setup entirely and provides the final arithmetic chain directly: “Sum = $10/2 * (2 + 20) = 110$. ##### 110”. A similar transition is observed in the GSM8k example: at step 30, the model walks through glass pricing logic with Python-like code and a `print()` statement for validation; by step 60, this is reduced to three compact sentences with no code, commentary, or explanation.

758

759

760

761

762

763

764

These early changes follow a consistent pattern. Narrative framing—phrases like “Let’s solve this by...” or “Now compute...” — are dropped first. Embedded symbolic or code-based verification disappears shortly after. While the arithmetic logic is preserved, the accompanying verbal scaffolding is eliminated. This compression begins after the d-schedule’s early burn-in phase (typically around step 20–40), as the model starts receiving stronger training signals to minimize token count while retaining correctness. By step 60, the pressure from the length penalty is sufficiently strong to shape model behavior measurably.

765

766

767

In summary, the model begins transitioning to a compressed reasoning style as early as step 60. The resulting outputs are more efficient and better aligned with inference-time brevity goals, but less transparent in their reasoning. This supports the hypothesis that the d-schedule enforces compression gradually,

with tangible effects emerging soon after step 30. While answer correctness remains intact, the interpretability of the solution path is the first to be sacrificed.

D.2 How Distillation Improves Reasoning Quality

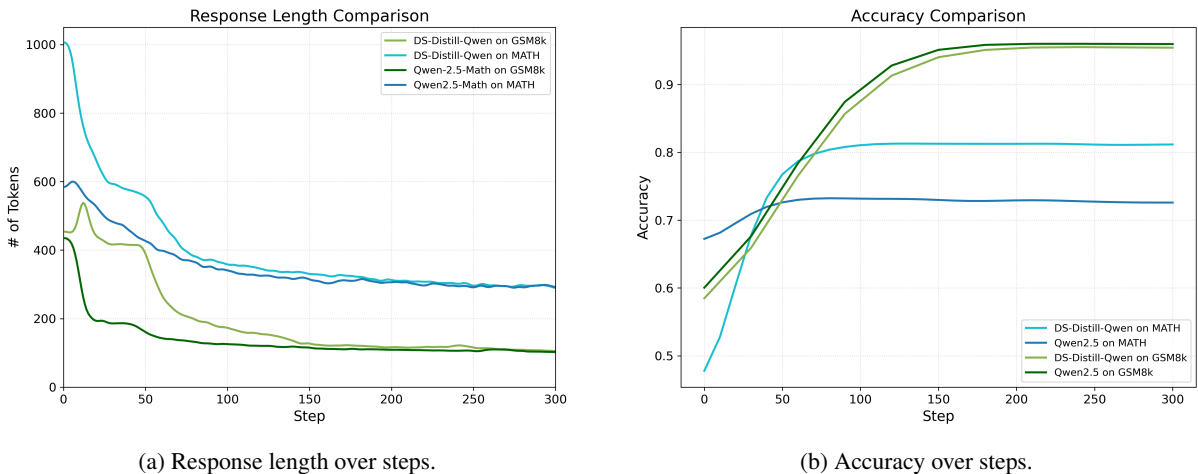


Figure 8: (a) Comparison of response lengths over training steps between DeepSeek-Distill-Qwen and Qwen2.5 models on GSM8k and MATH benchmarks. DeepSeek-Distill-Qwen consistently generates longer outputs, especially in earlier training phases. (b) Accuracy comparison over training steps between DeepSeek-Distill-Qwen and Qwen2.5 on GSM8k and MATH datasets. DeepSeek-Distill-Qwen achieves higher accuracy on both benchmarks. All model outputs are capped at 1000 tokens for fair comparison.

We applied AALC on both Deepseek-R1-Distill-Qwen and its original base Qwen2.5-Math. As shown in Figure 8a, on the GSM8k and MATH datasets, although the initial output length of Deepseek-R1-Distill-Qwen is greater than that of Qwen2.5-Math, the outputs of both models converge to the same length. In Figure 8b, on the MATH dataset, the accuracy of Deepseek-R1-Distill-Qwen is lower than that of Qwen2.5-Math due to its earlier output truncation. However, after fine-tuning with AALC, the accuracies of both Deepseek-R1-Distill-Qwen- and Qwen2.5-Math exhibit a steady increase.

We also analyze the qualitative differences. As shown in Table 7, the distilled model not only corrects the conceptual error in percentage conversion but also delivers a clearer and more structured explanation of the underlying computation. While the correction itself may reflect improvements in factual accuracy, the cleaner numerical logic and removal of redundant commentary suggest that structural refinement plays a supporting role. Thus, we believe this example offers evidence that performance gains under length-penalty and distillation can arise from improved reasoning organization, not just from truncation alone.

Dataset / Question	Base Model Output (Qwen2.5-Math-7B)	Distilled Model Output (Deepseek-R1-Distill-Qwen-7B)
What is the positive difference between 120% of 30 and 130% of 20?	Incorrect due to misinterpreting percentages: "120% of 30 is $0.12 \times 30 = 3.6$. 130% of 20 is $0.13 \times 20 = 2.6$. The positive difference is $3.6 - 2.6 = 1$. 1 "	Correct and concise: "120% of 30 is $1.2 \times 30 = 36$. 130% of 20 is $1.3 \times 20 = 26$. The difference is $36 - 26 = \boxed{10}$."

Table 7: Comparison of reasoning behavior between the base and distilled models after applying AALC. The base model misinterprets percentage values due to incorrect decimal logic and lacks structural clarity. The distilled model applies correct numeric conversions, removes redundant phrasing, and presents a concise, high-confidence reasoning chain—illustrating how compression improves both accuracy and readability.