



000
001
002
003
004
005
006
007
EFFICIENT REASONING WITH BALANCED THINKING
008009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
Anonymous authors
Paper under double-blind review
ABSTRACT
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
Large Reasoning Models (LRMs) have shown remarkable reasoning capabilities, yet they often suffer from overthinking, expending redundant computational steps on simple problems, or underthinking, failing to explore sufficient reasoning paths despite inherent capabilities. These issues lead to inefficiencies and potential inaccuracies, limiting practical deployment in resource-constrained settings. Existing methods to mitigate overthinking, such as suppressing reflective keywords or adjusting reasoning length, may inadvertently induce underthinking, compromising accuracy. Therefore, we propose REBALANCE, a training-free framework that achieves efficient reasoning with balanced thinking. REBALANCE leverages confidence as a continuous indicator of reasoning dynamics, identifying overthinking through high confidence variance and underthinking via consistent overconfidence. By aggregating hidden states from a small-scale dataset into reasoning mode prototypes, we compute a steering vector to guide LRMs' reasoning trajectories. A dynamic control function modulates this vector's strength and direction based on real-time confidence, pruning redundancy during overthinking, and promoting exploration during underthinking. Extensive experiments conducted on four models ranging from 0.5B to 32B, and across nine benchmarks in math reasoning, general question answering, and coding tasks demonstrate that REBALANCE effectively reduces output redundancy while improving accuracy, offering a general, training-free, and plug-and-play strategy for efficient and robust LRM deployment. Code and models will be made publicly available.

1 INTRODUCTION

032
033
034
035
036
037
038
Recent advances in Supervised Fine-Tuning (SFT) and Reinforcement Learning (RL) have substantially enhanced the reasoning capabilities of Large Reasoning Models (LRMs) (Jaech et al., 2024; Guo et al., 2025; Team, 2025). However, LRMs may exhibit *overthinking* (Chen et al., 2024b), allocating redundant reasoning steps to simple problems. This redundancy incurs substantial computational costs with marginal performance gains (Sui et al., 2025), and may introduce hallucinations (Sun et al., 2025). Thus, overthinking severely limits the practical deployment of LRMs in resource-constrained environments.039
040
041
042
043
044
045
046
047
048
049
050
Recent efforts (Yue et al., 2025) have been made to mitigate overthinking by shortening reasoning chains. However, these approaches primarily target overthinking and may overlook the critical issue of *underthinking* (Wang et al., 2025c), where LRMs fail to sufficiently explore valid reasoning paths despite possessing the inherent capability to solve the problem, as shown in Fig. 1(a). Specifically, Wang et al. (2025a), Ma et al. (2025b), and Chen et al. (2025) suppress keywords indicative of reflection and exploration, but indiscriminately affect both redundant and valuable reasoning, inevitably causing underthinking. Another direction (Zhang et al., 2025c; Lou et al., 2025; Huang et al., 2025a) adjusts reasoning length based on problem difficulty via SFT or RL, yet often penalizes lengthy reasoning (Su et al., 2025b) or dilutes rewards for control tokens (Fang et al., 2025). Such designs may cause decision boundary collapse (Lou et al., 2025), biasing models toward overly short reasoning chains and inducing underthinking. Hence, a key question arises: *How can we mitigate overthinking without inducing underthinking, achieving efficient reasoning with balanced thinking?*051
052
053
Key observations. To address this issue, we need to develop a dynamic mechanism capable of explicitly modeling and controlling both overthinking and underthinking. Though recent works (Zhang et al., 2025a; Yang et al., 2025b; Lin et al., 2025a) have achieved dynamic control by adopting manually designed metrics to adaptively retain or discard entire reasoning paths, this rigid binary selection



(a) Qualitative comparison.

(b) Quantitative comparison.

Figure 1: Qualitative and quantitative comparisons with previous state-of-the-art methods for mitigating overthinking. (a) Given the question “For what real values of x is $-4 < x^4 + 4x^2 < 21$?", the model first obtains intervals $(-\sqrt{3}, 0)$ and $(0, \sqrt{3})$, and then verifies if $x = 0$ is included. However, the baseline (Guo et al., 2025) redundantly checks irrelevant values after correctly validating $x = 0$, causing overthinking. Current mitigation methods (Yang et al., 2025b) overly suppress necessary reflection, leading to underthinking. Our method dynamically controls the reasoning state, effectively balancing these two extremes. (b) REBALANCE outperforms previous state-of-the-art method (Chen et al., 2025) across multiple mathematical reasoning datasets and model scales (0.5B–32B), reducing reasoning length while simultaneously improving accuracy.

may sacrifice the potentially valuable intermediate reasoning steps, thus still risking underthinking. This motivates us to investigate a continuous and reliable indicator of reasoning states for providing dynamic fine-grained reasoning control.

As shown in Fig. 2, we can observe that the confidence values correlate with LRM’s reasoning behaviors. Specifically, high confidence variance may reflect frequent indecisive switching between different reasoning paths, causing redundant steps and delayed answer convergence, *i.e.*, *overthinking*. Conversely, consistent overconfidence can lead to premature commitment to incorrect reasoning paths, *i.e.*, *underthinking*. Thus, confidence can be leveraged as an indicator of reasoning dynamics. Given that LRM’s internal reasoning states are inherently represented by their hidden states (Su et al., 2025a), this observation prompts us to consider *whether the efficient reasoning can be achieved through balanced thinking, by dynamically adjusting hidden states according to confidence levels*.

Our solution. In this work, we propose **ReBalance**, a training-free method that achieves efficient Reasoning with **Balanced** thinking. To achieve dynamic control between overthinking and underthinking, we first identify reasoning steps indicating overthinking and underthinking from a small-scale seen dataset, aggregate their corresponding hidden states into reasoning mode prototypes, and compute a steering vector that encodes the transition between them, *i.e.*, from overthinking to underthinking. Since the steering vector captures the model’s inherent reasoning dynamics, it exhibits strong generalization across diverse unseen data, as demonstrated in our experiments.

With this steering vector, we further introduce a dynamic control function that modulates the strength and direction of the vector based on the model’s confidence at each step. When signs of overthinking emerge, the steering is amplified to prune redundancy. Conversely, when underthinking is inferred, steering is reversed to promote exploration of alternative reasoning paths. This adaptive mechanism effectively balances reasoning depth across various contexts, enhancing efficiency without compromising the core reasoning abilities.

Extensive experiments across four models ranging from 0.5B to 32B, and on nine benchmarks covering math reasoning, general question answering, and coding tasks, demonstrate the effectiveness and strong generalization capabilities of REBALANCE. Notably, REBALANCE not only reduces output length but also improves the accuracy. To summarize, our contributions are as follows:

- As the current methods struggle to balance between overthinking and underthinking, we identify that confidence can serve as a continuous and reliable signal for characterizing both overthinking and underthinking in LRM’s, enabling fine-grained behavioral control.

- 108 • To achieve dynamic reasoning control, we propose REBALANCE, an efficient and training-
109 free framework that dynamically steers the reasoning trajectory of LRM_s by modulating
110 their internal state based on confidence estimates.
- 111 • Extensive experiments across different models and tasks demonstrate that REBALANCE
112 improves both inference efficiency and accuracy, offering a plug-and-play solution for
113 boosting the efficiency of LRM_s without compromising performance.

115 2 BACKGROUND AND MOTIVATION

116 2.1 PRELIMINARIES

117 In the following, to investigate the dynamics of the reasoning process of large reasoning models
118 (LRM_s), we introduce the computation of *stepwise confidence* and *confidence variance*. Stepwise
119 confidence measures the degree to which the model consistently adheres to the same reasoning path,
120 while confidence variance between different steps quantifies the frequency of switching between
121 different reasoning paths. The discussion of related work is presented in Appendix G.

122 **Stepwise confidence.** For each token position $t \in \mathcal{T}_s$, we can define the tokenwise maximum
123 predicted probability $p_t^{\max} = \max_{v \in V} \mathbf{p}_\theta(v | x_{<t})$. Then, we can obtain the confidence c_s of the
124 reasoning step S_s , which is the geometric average of these maxima across all tokens in the step:

$$125 c_s = \exp \left(\frac{1}{|\mathcal{T}_s|} \sum_{t \in \mathcal{T}_s} \ln p_t^{\max} \right) \quad (1)$$

126 **Confidence variance.** To capture short-term fluctuations in confidence, we compute the confi-
127 dence variance $\text{Var}(\cdot)$ over recent steps. Since long-term history is less relevant, we focus on local
128 variability by calculating the variance within a sliding window of size $|W| \geq 1$, and we can define
129 the window \mathcal{W}_s for the s -th step as $\mathcal{W}_s = \{\max(1, s-W+1), \dots, s\}$. Then, with the average step
130 confidence $\bar{c}_s = \frac{1}{|\mathcal{W}_s|} \sum_{j \in \mathcal{W}_s} c_j$ within the window \mathcal{W}_s , we can obtain the confidence variance for
131 the s -th step $\text{Var}(c_s; \mathcal{W}_s)$ as:

$$132 \text{Var}(c_s; \mathcal{W}_s) = \begin{cases} 0, & |\mathcal{W}_s| = 1, \\ \frac{1}{|\mathcal{W}_s|} \sum_{j \in \mathcal{W}_s} (c_j - \bar{c}_s)^2, & |\mathcal{W}_s| \geq 2. \end{cases} \quad (2)$$

133 To this end, regardless of the current stepwise confidence level, a high $\text{Var}(c_s; \mathcal{W}_s)$ indicates fre-
134 quent switching among different reasoning paths, which may force the model to continue generating
135 redundant reasoning steps instead of concluding, leading to *overthinking*. Differently, consistently
136 high c_s with low $\text{Var}(c_s; \mathcal{W}_s)$ implies premature commitment and potential *underthinking*. These
137 statistics will guide the dynamic control mechanism that will be introduced later.

138 2.2 KEY OBSERVATIONS

139 As discussed above, existing approaches designed to mitigate overthinking effectively reduce the
140 length of inference outputs, yet struggle to achieve satisfactory accuracy. To investigate the underly-
141 ing reasons, we analyze how the length of reasoning sequences relates to the ground-truth reasoning
142 length for both correctly and incorrectly answered samples, before and after applying methods in-
143 tended to mitigate overthinking, as shown in Fig. 2(a). Specifically, we collect inference samples
144 under three conditions: the original model, the model after applying existing methods, and the model
145 after applying our proposed method. We utilize ground-truth as a proxy for ideal reasoning length.

146 **The trade-off between overthinking and underthinking.** Theoretically, if an overthinking miti-
147 gation approach effectively reduces redundant reasoning steps, the reasoning sequence lengths of
148 correctly answered samples should accordingly decrease. Conversely, if such methods introduce un-
149 derthinking by prematurely truncating necessary reasoning, resulting in errors, the reasoning lengths
150 for these incorrect samples should also decrease. As shown in Fig. 2(a), both existing methods and




Figure 2: (a) Effects of overthinking mitigation on reasoning modes. We compare the distributions of reasoning lengths for correct and incorrect predictions before and after applying overthinking mitigation methods. The reduction in reasoning lengths for correct and incorrect predictions indicates the degree to which overthinking is alleviated and underthinking is introduced, respectively. Existing methods significantly introduce underthinking, whereas our method effectively achieves a balanced reduction of both. **(b) Correlation between confidence and reasoning modes.** We observe that the overthinking samples exhibit higher confidence variance compared to normal samples, while underthinking samples show persistently high confidence levels.

our proposed approach significantly mitigate overthinking. However, existing methods introduce notable underthinking, whereas our proposed approach maintains reasoning length distribution similar to the original model, demonstrating superior balanced thinking capacity.

Consequently, addressing the critical issue of simultaneously mitigating overthinking and preventing underthinking becomes essential. Achieving this requires explicit modeling of these two reasoning modes. Intuitively, questions correctly answered by the original model but incorrectly answered after applying overthinking mitigation methods are likely due to restricted exploration, indicating underthinking. Conversely, questions correctly answered by both the original and mitigated models with shortened reasoning sequences likely reflect the successful reduction of redundant steps, indicating overthinking. Based on these categorizations, we analyze changes in stepwise confidence and confidence variance relative to normal reasoning, as illustrated in Fig. 2(b).

Confidence indicates reasoning states. Our analysis reveals that overthinking typically coincides with higher confidence variance, indicative of hesitation across reasoning steps, while underthinking is characterized by persistently high confidence levels, reflecting premature commitment to incorrect reasoning paths without sufficient exploration. These findings support our proposal that confidence can serve as a continuous and reliable indicator of the model’s reasoning state, enabling fine-grained behavioral control. A comprehensive analysis, including the correlation between confidence and reasoning length (Appendix A.2), inertia effects of confidence states (Appendix A.3), confidence variations across models (Appendix A.4), model keywords and confidence states (Appendix A.6), and the discriminability of confidence in latent space (Appendix A.5) are provided in the Appendix.

3 METHOD

3.1 OVERVIEW

In this section, we present REBALANCE, a training-free framework designed to dynamically balance overthinking and underthinking, thereby improving efficiency without compromising accuracy.

Specifically, REBALANCE first explicitly models reasoning states prone to overthinking or underthinking using stepwise confidence and confidence variance (Sec. 3.2). Next, it utilizes these identified states to extract distinct steering vectors from deep-layer hidden states, capturing key behavioral patterns of different reasoning modes between overthinking and underthinking (Sec. 3.3). Finally, the steering vectors will be controlled by a dynamic function that adaptively modulates steering strength and direction, ensuring balanced thinking during the reasoning process (Sec. 3.4). Collectively, these complementary components provide precise, adaptive, and efficient control over the reasoning process. The overview is shown in Fig. 3.




Figure 3: Illustration of the REBALANCE framework. We first perform offline one-pass data collection on a small-scale seen dataset. At each step, the steering vector is extracted at the first token of the specified layer based on confidence, and a dynamic function is fitted according to model behaviors. During deployment, the dynamic function outputs steering weights based on the model’s real-time confidence online, thus balancing between overthinking and underthinking.

3.2 EXPLICIT MODELING OF OVERTHINKING AND UNDERTHINKING

Building upon the insights that confidence serves as a reliable indicator of overthinking and underthinking, we first formally define these reasoning states and then explicitly model them using confidence metrics.

Definitions of overthinking and underthinking. Let the $\langle \text{think} \rangle \dots \langle / \text{think} \rangle$ trajectory be segmented into steps S_1, \dots, S_{\max} by the delimiter mentioned in Sec. 2.1. Denote the partial reasoning up to step s by $r_{\leq s}$ and the induced answer distribution (if forced to stop at s) by π_s ; let prediction $d_s = \arg \max \pi_s$ under a specified decoding rule, then we define the *stability index* as:

$$s^* = \min \{s : d_{s'} = d_s \text{ for all } s' \geq s \text{ and } d_s \text{ is correct}\}. \quad (3)$$

The stability index s^* serves as a signal to distinguish different reasoning modes. Specifically, A trajectory may exhibit *overthinking* if it continues after s^* . Conversely, it exhibits *underthinking* if it stops at step s with incorrect prediction d_s while there exists $s' > s$ with correct $d_{s'}$. These definitions formalize the notions of redundant computation after convergence to the correct answer and premature termination before sufficient reasoning.

Explicit modeling with confidence. Then, the above definitions can be instantiated using the step-wise confidence c_s and the confidence variance $v_s = \text{Var}(c_s; \mathcal{W}_s)$ introduced in Sec. 2.1. With a small-scale seen dataset that has been used for training, we can obtain the empirical quantiles (Hyndman & Fan, 1996) $Q_c(\cdot)$ and $Q_v(\cdot)$ and thresholds as:

$$\tau_c^L = Q_c(q_L), \quad \tau_c^H = Q_c(q_H), \quad \tau_v^L = Q_v(q_L), \quad \tau_v^H = Q_v(q_H), \quad (4)$$

where $0 < q_L < q_H < 1$ specify the lower and upper quantiles, respectively. Then, with these thresholds, we can classify the reasoning steps into two sets \mathcal{O} and \mathcal{U} :

$$\mathcal{O} \leftarrow \{s : c_s \leq \tau_c^L \wedge v_s \geq \tau_v^H\}, \quad \mathcal{U} \leftarrow \{s : c_s \geq \tau_c^H \wedge v_s \leq \tau_v^L\}. \quad (5)$$

Concretely, as illustrated in Fig. 2(b), the overthinking set \mathcal{O} contains instances characterized by high reasoning variance and low confidence, reflecting unstable or oscillating reasoning trajectories. On the other hand, the underthinking set \mathcal{U} comprises cases with low variance and persistently high confidence, indicating premature convergence and a tendency toward underthinking. Instances not belonging to $\mathcal{O} \cup \mathcal{U}$ can be treated as *normal* and are excluded from further analysis.

270 3.3 CONFIDENCE-BASED STEERING VECTOR EXTRACTION
271

272 In this section, based on the modeling of overthinking and underthinking introduced in Sec 3.2, we
273 extract prototypical representations of both reasoning modes from the hidden states of LRM_s via
274 an offline, single forward pass. Then, the resulting prototypes enable the construction of a steering
275 vector that delineates the trajectory from overthinking to underthinking, thereby facilitating fine-
276 grained behavior control.

277 **One-pass prototype extraction.** To obtain prototypes, we perform a single offline inference pass
278 over a small seen dataset $\mathcal{D}_{\text{seen}}$, segmenting reasoning steps by the delimiter $\backslash n \nbackslash n$. During this pass,
279 we automatically select the optimal deep-layer based on LRM_s' intrinsic separability of reasoning
280 modes (see Appendix A.5), from which we collect hidden states $\mathbf{h}_{t_s^{(1)}}$ at the first token $t_s^{(1)}$ of each
281 step. $\mathbf{h}_{t_s^{(1)}}$ serves as a compact encoding of step-level intent (Yang et al., 2025b) and, under causal
282 masking, conditions the generation of all subsequent tokens within the step. We find that deeper layers
283 exhibit stronger discriminability between reasoning modes and improved generalization across
284 datasets, as analyzed in Appendix A.5.
285

286 Then, with the hidden stages and the tags \mathcal{O} and \mathcal{U} mentioned in Sec. 3.2 for each step, we can
287 obtain the overthinking and underthinking prototypes, *i.e.*, μ^{O} and μ^{U} , respectively:

$$288 \quad \mu^{\text{O}} = \frac{1}{|\mathcal{O}|} \sum_{s \in \mathcal{O}} \mathbf{h}_{t_s^{(1)}}, \quad \mu^{\text{U}} = \frac{1}{|\mathcal{U}|} \sum_{s \in \mathcal{U}} \mathbf{h}_{t_s^{(1)}}. \quad (6)$$

291 **Steering vector construction.** The prototypes μ^{O} and μ^{U} denote the representations leading to
292 overthinking and underthinking respectively. The steering vector is then defined as the direction
293 from underthinking μ^{U} to overthinking μ^{O} :

$$294 \quad \mathbf{v} = \frac{\mu^{\text{O}} - \mu^{\text{U}}}{\|\mu^{\text{O}} - \mu^{\text{U}}\|_2}. \quad (7)$$

295 With the steering vector \mathbf{v} , we can formalize the transition between two reasoning modes. To mod-
296 ule the behavior during inference, we adjust the initial token $\mathbf{h}_{t_s^{(1)}}$ of each step as follows:

$$297 \quad \tilde{\mathbf{h}}_{t_s^{(1)}} = \mathbf{h}_{t_s^{(1)}} + \alpha_s \mathbf{v}, \quad \alpha_s = \lambda_s \delta_s, \quad \lambda_s \geq 0, \quad \delta_s \in \{+1, -1\}, \quad (8)$$

301 where α_s represents the signed steering weight at step s , combining the steering strength λ_s and
302 direction δ_s . When $\delta_s = +1$, we can address underthinking by stimulating the exploration of alter-
303 native reasoning paths. Conversely, $\delta_s = -1$ mitigates overthinking by encouraging commitment.
304 These adjustments conceptually establish the boundaries within which the model's reasoning pro-
305 cess operates, aiming to maintain a balanced state that ensures efficient and effective reasoning.

306 3.4 MODEL BEHAVIOR-BASED DYNAMIC CONTROL FUNCTION
307

308 Considering the evolving nature of model states and contexts over time, we introduce a dynamic
309 control function that adaptively adjusts steering strength and direction during inference. Motivated
310 by Sec. 2.2, which shows that the confidence correlates with reasoning modes, the steering weight α_s
311 can be deemed as the output of a continuous function $g(c_s, v_s)$ with respect to the current confidence
312 c_s and variance v_s . Therefore, the steering weight α_s , strength λ_s and direction δ_s are defined as:

$$313 \quad \alpha_s = g(c_s, v_s) = \delta_s \cdot \lambda_s. \quad (9)$$

314 During inference, at each step s , we obtain the confidence c_s and variance v_s , set $\alpha_s = g(c_s, v_s)$, and
315 inject $\alpha_s \mathbf{v}$ at the first token $t_s^{(1)}$ for the selected layer as in Eq. 8. This keeps trajectories between the
316 overthinking and underthinking boundaries while adding no extra forward passes beyond standard
317 decoding. Concretely, the dynamic control function $g(c_s, v_s)$ formulates as:

$$319 \quad g(c_s, v_s) = \underbrace{\text{sign}(c_s - \tau_c^{\text{H}})}_{\text{Steering direction } \delta_s} \cdot \underbrace{B(c_s, v_s) \tanh(|c_s - \tau_c^{\text{H}}|)}_{\text{Steering strength } \lambda_s} \quad (10)$$

322 **The steering direction** δ_s is determined by the sign function $\text{sign}(c_s - \tau_c^{\text{H}})$, where the confidence
323 threshold τ_c^{H} is obtained as Eq. 4. It takes a negative value when confidence is below the high-
324 confidence threshold ($c_s < \tau_c^{\text{H}}$) to mitigate overthinking, and a positive value when confidence is

324	325	MATH-500		AIME24		AIME25		GSMBK		AMC23		Olympiad	
		Method	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑
DeepSeek-R1-Distill-Qwen-1.5B													
326	Baseline (Guo et al., 2025)	79.6	4516	23.3	12596	16.7	14556	76.0	1018	55.0	8990	41.2	8785
327	CoD (Xu et al., 2025a)	80.2	3512	33.3	11894	13.3	10941	69.5	531	62.5	6812	35.2	7160
328	DEER (Yang et al., 2025b)	67.0	2251	20.0	8135	23.3	8719	69.2	684	57.5	5132	35.4	5982
329	NoThinking (Ma et al., 2025b)	75.0	1582	6.67	6998	16.6	7473	61.6	285	60.0	3267	37.3	3485
330	NoWait (Wang et al., 2025a)	78.0	2645	30.0	8225	16.6	7574	75.1	641	60.0	3302	40.6	4795
331	Dynasor-CoT (Fu et al., 2025)	77.2	3694	26.7	10564	26.7	12462	77.1	1035	72.5	6505	42.6	8859
332	SEAL (Chen et al., 2025)	78.6	3259	23.3	10785	26.7	8544	76.4	754	75.0	5084	32.7	7117
333	Manifold Steering (Huang et al., 2025b)	78.6	3458	30.0	10134	—	—	77.2	593	72.5	6236	—	—
334	ReBalance (Ours)	83.0	3474	36.7	9040	30.0	8140	78.3	765	80.0	5216	43.9	7235
335	△ vs. Baseline	(+3.4)	(-23.1%)	(+13.4)	(-28.2%)	(+13.3)	(-44.1%)	(+2.3)	(-24.9%)	(+25.0)	(-42.0%)	(+2.7)	(-17.6%)
DeepSeek-R1-Distill-Qwen-7B													
336	Baseline (Guo et al., 2025)	89.8	3699	40.0	13994	26.7	13778	89.2	1098	75.0	6898	56.1	7590
337	CoD (Xu et al., 2025a)	90.0	3127	46.7	11663	36.7	10198	84.5	339	85.0	3654	47.5	5688
338	DEER (Yang et al., 2025b)	87.8	2367	50.0	8924	40.0	8919	90.4	676	80.0	5157	53.9	5804
339	NoThinking (Ma et al., 2025b)	80.6	834	26.7	4427	20.0	7850	87.1	284	65.0	1911	45.3	3331
340	NoWait (Wang et al., 2025a)	86.8	2479	50.0	6844	26.7	6979	90.2	806	85.0	3795	52.1	4760
341	Dynasor-CoT (Fu et al., 2025)	88.2	2723	46.7	9864	33.3	11069	87.6	732	85.0	5121	55.4	7427
342	SEAL (Chen et al., 2025)	90.6	2843	43.3	10112	26.7	9835	88.4	811	77.5	5164	53.9	6261
343	Manifold Steering (Huang et al., 2025b)	88.4	2239	53.3	8457	—	—	87.6	440	87.5	4440	—	—
344	ReBalance (Ours)	92.6	2903	56.7	9012	40.0	9227	91.6	912	95.0	4767	57.0	6321
345	△ vs. Baseline	(+2.8)	(-21.5%)	(+10.0)	(-19.7%)	(+13.3)	(-33.0%)	(+2.4)	(-16.9%)	(+20.0)	(-30.9%)	(+0.9)	(-16.2%)
Qwen3-14B													
346	Baseline (Yang et al., 2025a)	93.8	4470	66.7	10888	56.7	13125	95.1	2231	95.0	7240	60.6	7450
347	CoD (Xu et al., 2025a)	93.8	2950	66.7	10212	53.3	11828	95.6	627	95.0	5360	62.2	6554
348	DEER (Yang et al., 2025b)	93.0	2825	66.7	9973	56.7	11806	95.8	934	95.0	5527	66.1	6849
349	NoThinking (Ma et al., 2025b)	93.8	2657	70.0	8898	53.3	9892	95.1	369	87.5	4503	64.0	5880
350	NoWait (Wang et al., 2025a)	92.8	3219	60.0	10507	56.7	10924	95.6	1129	95.0	5050	59.2	7332
351	Dynasor-CoT (Fu et al., 2025)	93.8	4063	73.3	10369	60.0	12159	95.6	1483	95.0	6582	—	—
352	SEAL (Chen et al., 2025)	93.4	3727	63.3	10322	50.0	10901	95.7	1369	90.0	6126	62.3	7131
353	ReBalance (Ours)	94.0	3641	73.3	9464	56.7	11057	96.3	1441	100.0	5230	66.3	7257
354	△ vs. Baseline	(+0.2)	(-18.5%)	(+6.6)	(-13.1%)	(+0.0)	(-15.8%)	(+1.2)	(-35.4%)	(+5.0)	(-27.8%)	(+5.7)	(-2.6%)
QwQ-32B													
355	Baseline (Teun, 2025)	94.8	4535	66.7	14342	46.7	13350	96.3	1506	87.5	7021	66.7	8219
356	CoD (Xu et al., 2025a)	93.8	3516	63.3	11438	46.7	12189	96.2	670	92.5	6217	67.7	7028
357	DEER (Yang et al., 2025b)	94.4	3179	70.0	8885	46.7	10972	96.2	944	95.0	6435	64.3	7085
358	NoThinking (Ma et al., 2025b)	94.8	3912	66.7	10507	56.7	11839	96.5	1326	90.0	7119	66.1	8132
359	NoWait (Wang et al., 2025a)	93.8	2879	66.7	8190	63.3	8970	96.3	942	92.5	4717	62.6	8223
360	Dynasor-CoT (Fu et al., 2025)	94.2	4176	63.3	11156	—	—	95.2	1095	90.0	6544	—	—
361	SEAL (Chen et al., 2025)	92.6	3536	63.3	10344	56.7	11384	96.2	1221	95.0	6341	67.5	7371
362	FlashThink (Jiang et al., 2025)	93.2	3144	60.0	10034	40.0	11861	96.5	910	92.5	6702	—	—
363	TrimR (Lin et al., 2025a)	93.8	3830	56.7	8345	43.3	8827	93.7	1319	90.0	6055	—	—
364	ReBalance (Ours)	95.2	3662	70.0	10350	63.3	11575	96.8	1289	95.0	6064	68.6	7422
365	△ vs. Baseline	(+0.4)	(-19.3%)	(+3.3)	(-27.8%)	(+16.6)	(-13.3%)	(+0.5)	(-14.4%)	(+7.5)	(-13.6%)	(+1.9)	(-9.7%)

Table 1: Performance on math reasoning benchmarks. Metrics include Pass@1 (↑) and #Tokens (↓) on six math reasoning benchmarks. Changes are shown in orange for Pass@1 and blue for #Tokens. FlashThink and TrimR are reproduced according to the paper.

above this threshold ($c_s > \tau_c^H$) to alleviate underthinking. This guarantees the steering consistently directs the state away from the nearer reasoning boundary.

The steering strength λ_s is composed of two parts: (1) *soft saturation* $\tanh(|c_s - \tau_c^H|)$ and (2) *variance-aware amplitude* $B(c_s, v_s)$. Specifically, regarding the soft saturation $\tanh(|c_s - \tau_c^H|)$, a smooth, saturating growth in $|c_s - \tau_c^H|$ avoids abrupt changes and keeps the mapping monotone in c_s for any fixed v_s . The soft saturation function guarantees the steering strength grows gradually as the state approaches the reasoning boundary, ensuring numerical stability.

Differently, the variance-aware amplitude $B(c_s, v_s)$ is a model behavior-based scalar amplitude that adapts across models based on the step confidence c_s and variance v_s . It is required to indicate the model’s current thinking status, shifting between moderate and overthinking/underthinking reasoning modes. To this end, the amplitude function can be formulated as:

$$B(c_s, v_s) = \begin{cases} B_m + (B_o - B_m)\psi(c_s, v_s) & \text{if } c_s \leq \tau_c^L \text{ and } v_s \geq \tau_v^H, \\ B_m + (B_u - B_m)\psi(c_s, v_s) & \text{if } c_s \geq \tau_c^H \text{ and } v_s \leq \tau_v^L, \\ B_m & \text{otherwise.} \end{cases} \quad (11)$$

In Eq. 11, B_m , B_o , and B_u are adaptive mode boundaries representing moderate, overthinking, and underthinking, respectively. $\psi(c_s, v_s)$ denotes a conditioned gating function whose output ranges from 0 to 1 to ensure smooth transitions. The thresholds $(\tau_c^L, \tau_c^H, \tau_v^L, \tau_v^H)$ are obtained in Eq. 4. Following the reasoning mode definitions outlined in Eq. 5, when $c_s \leq \tau_c^L$ and $v_s \geq \tau_v^H$, indicating a state of overthinking, the transition occurs between B_m and B_o . Differently, when $c_s \geq \tau_c^H$ and $v_s \leq \tau_v^L$, indicating a state of underthinking, the transition should be performed between B_m and B_u . Notably, B_m and B_o are adaptively derived from models without manual tuning.

In this context, the amplitude $B(c_s, v_s)$ serves as an indicator of the current reasoning status, complemented by the saturation function, which ensures the numerical stability of the final steering strength. More details, theoretical derivations, and proofs regarding the mode boundaries and the gating function are provided in Appendix B due to the page limit.

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392

393

394

395

396

397

398

399

400

401

402

403

404

405

406

407

408

Method	SCIENCE		COMMONSENSE		PROGRAMMING	
	Pass@1↑	GPQA-D #Tokens↓	Pass@1↑	StrategyQA #Tokens↓	Pass@1↑	LiveCodeBench #Tokens↓
DeepSeek-R1-Distill-Qwen-1.5B						
Baseline	17.1	8 727	63.2	435	19.5	12 509
Ours	21.7 (+4.6)	6 902 (-20.9%)	67.7 (+4.5)	401 (-7.8%)	22.5 (+3.0)	11 622 (-7.1%)
DeepSeek-R1-Distill-Qwen-7B						
Baseline	33.8	7 392	88.1	350	44.0	9 851
Ours	39.4 (+5.6)	5 180 (-29.9%)	88.9 (+0.8)	310 (-11.4%)	46.5 (+2.5)	8 651 (-12.2%)
Qwen3-14B						
Baseline	60.6	7 451	94.2	267	83.5	7 101
Ours	67.2 (+6.6)	5 779 (-22.4%)	94.3 (+0.1)	260 (-2.6%)	84.6 (+1.1)	6 088 (-14.3%)
QwQ-32B						
Baseline	63.1	7 424	93.6	274	87.5	6 622
Ours	67.2 (+4.1)	6 296 (-15.2%)	95.7 (+2.1)	265 (-3.3%)	88.3 (+0.8)	5 649 (-14.7%)

Table 2: Generalization capabilities on other non-math tasks. Metrics include Pass@1 (↑) and #Tokens (↓). Changes are shown in orange for Pass@1 and blue for #Tokens.

Method	Math500		GSM8K		Olympiad	
	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓
Rebalance (Ours)	83.0	3474	78.3	765	43.9	7235
$ \mathcal{W}_s = 5$	82.4	3686	78.2	910	42.2	7343
$ \mathcal{W}_s = 10$	81.0	4084	77.8	940	44.7	7679
$B_u = 0$	82.0	3343	78.1	761	39.4	7147
$B_u = 0.2$	83.6	3600	80.5	850	44.2	7923
$B_u = 0.5$	81.4	3543	77.7	890	41.3	8773

Table 3: Ablations on the R1-1.5B backbone across difficulty levels. We analyze performance changes on three math benchmarks with varying difficulty: **Math500** (medium), **GSM8K** (easy), and **Olympiad** (hard). Metrics are Pass@1 accuracy (%) and generated token numbers. Arrows indicate change relative to our original REBALANCE: Acc. ↑ increase, ↓ decrease; Tokens ↓ decrease, ↑ increase.

4 EXPERIMENT

Evaluation is conducted on *mathematics reasoning* datasets: MATH-500 (Lightman et al., 2023b), AIME24 (AI-MO, 2024a), AIME25 (OpenCompass, 2025), AMC23 (AI-MO, 2024b), GSM8K (Cobbe et al., 2021), and OLYMPIADBENCH (He et al., 2024); *scientific reasoning* dataset, GPQA DIAMOND (Rein et al., 2024); *commonsense reasoning* dataset, STRATEGYQA (Geva et al., 2021); and *code reasoning* dataset, LIVECODEBENCH (Jain et al., 2024). Besides, the proposed steering extraction and dynamic control function fitting are performed for each backbone once and held fixed across all unseen benchmarks for evaluation. 500 randomly sampled MATH (Hendrycks et al., 2021) problems are utilized during these processes, and the sensitivity analysis is shown in Fig. 5(c). More comprehensive experimental details, [including the baseline introductions](#), are provided in Appendix D.

4.1 MAIN RESULTS

Math reasoning. As shown in Tab. 1, REBALANCE outperforms all baselines on six math reasoning benchmarks spanning diverse difficulties and distributions. Without introducing any auxiliary models or inference stages, it simultaneously improves Pass@1 and reduces the average generated token count by at most 52.3%. Notably, on AMC23, REBALANCE attains perfect Pass@1 with Qwen3-14B and it lifts DeepSeek-R1-Distill-Qwen-7B to performance comparable to QwQ-32B.

Other reasoning scenarios. We evaluate REBALANCE in a cross-domain setting, fixing the steering vector and control surface across tasks. As shown in Tab. 2, without domain-specific tuning, REBALANCE maintains Pass@1 and reduces the reasoning length by at most 47.5% on challenging scientific reasoning, programming, and simpler commonsense QA tasks. These results demonstrate strong cross-domain generalization.




Figure 5: (a-b) Layerwise Performance of MATH-500 for (a) R1-7B and (b) QwQ-32B. (c) Sensitivity to sample size for steering vector extraction. (d) Performance with cross-domain vectors.

4.2 ABLATION STUDY

Impact of static α_s control. We ablate the dynamic schedule by fixing the steering weight α_s . As shown in Fig. 4, positive α_s ($\mathcal{U} \rightarrow \mathcal{O}$) improves accuracy but increases reasoning length, intensifying with larger $|\alpha_s|$ (e.g., $\alpha_s = +3$ on QwQ-32B yields a 147% token increase). Negative α_s reduces tokens at the expense of accuracy. These results motivate dynamic adapting α_s to instance difficulty.

Impact of the steering layer. We test generalization by fitting *steering vectors* and *control surfaces* at various layers (selected by depth ratio). Fig. 5(a-b) shows that steering at any tested depth reduces token count without harming accuracy. The strongest trade-off appears in mid-to-late layers, consistent with our probing analysis (Appendix A.5), where representations from these layers exhibit the highest confidence separability and thus incur minimal noise.

Steering vector choice and generalization. As shown in Fig. 5(c-d), we estimate steering vectors from math datasets of varying scale and a cross-domain science corpus (GPQA), fitting a *control surface* for each. Vectors generalize across datasets: REBALANCE improves efficiency while preserving accuracy. A clear trend emerges: vectors from harder datasets prioritize accuracy gains over token savings, aligning with the method’s mechanism: harder data induces a *conservative* surface prioritizing correctness, while easier data yields a more *aggressive* one favoring token savings.

Impact of window size $|\mathcal{W}_s|$. Rows 4 and 5 of Tab. 9 show how the control-surface window size $|\mathcal{W}_s|$ affects reasoning. Empirically, larger windows substantially increase token usage, consistent with the intuition that they smooth short-term fluctuations while reducing responsiveness to local anomalies. Theoretically, we show that the confidence trajectory during inference satisfies a Markovian continuity assumption (see Appendix A.3); a small window ($|\mathcal{W}_s| = 2$) is therefore sufficiently expressive and more sensitive to local reasoning patterns. However, with a larger window, accuracy increases on the hard *Olympiad* benchmark, in line with prior findings that extended deliberation improves performance at the expense of longer outputs (Jin et al., 2024; Muennighoff et al., 2025).

Impact of underthinking mode boundary B_u . Tab. 9 (Row 6) shows that removing the underthinking mode boundary slightly reduces tokens but significantly lowers accuracy, especially on tasks demanding extended reasoning. Moderate increases ($0.1 \rightarrow 0.2$; Rows 7-8) encourage deeper deliberation and boost accuracy at a modest token cost, while larger increases ($0.1 \rightarrow 0.5$) trigger overthinking and degrade performance. An overly strong boundary disrupts these normal reasoning paths. The other two mode boundaries B_m and B_o , are adaptively determined based on model behavior, requiring no manual tuning (Appendix B).

5 CONCLUSION

This paper analyzes the limitations of existing approaches to overthinking mitigation, and we observe that such attempts often introduce the countervailing problem of underthinking. Therefore, we propose REBALANCE, a training-free method that curbs overthinking while avoiding underthinking. Extensive experiments across diverse models and datasets show that REBALANCE reduces redundancy while preserving accuracy, achieving efficient reasoning with balanced thinking. A promising future direction is to apply REBALANCE to the multi-modal reasoning scenarios.




Figure 4: Static α_a Control on MATH-500. Top: R1-7B; Bottom: QwQ-32B.

486 REFERENCES
487

488 Pranjal Aggarwal and Sean Welleck. L1: Controlling how long a reasoning model thinks with
489 reinforcement learning, 2025. *URL* <https://arxiv.org/abs/2503.04697>, 2025.

490 AI-MO. Aime 2024, July 2024a. *URL* <https://huggingface.co/datasets/AI-MO/aimo-validation-aime>.

491

492 AI-MO. Amc 2023, July 2024b. *URL* <https://huggingface.co/datasets/AI-MO/aimo-validation-amc>.

493

494 Anthropic. Claude 3.5 Sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>, 2024. Accessed: 2025-11-22.

495

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

498

499 Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gian-
500 inazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczek, et al. Graph of
501 thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI*
502 conference on artificial intelligence, volume 38, pp. 17682–17690, 2024.

503

504 Qiguang Chen, Libo Qin, Jiaqi Wang, Jingxuan Zhou, and Wanxiang Che. Unlocking the capa-
505 bilities of thought: A reasoning boundary framework to quantify and optimize chain-of-thought.
506 *Advances in Neural Information Processing Systems*, 37:54872–54904, 2024a.

507

508 Runjin Chen, Zhenyu Zhang, Junyuan Hong, Souvik Kundu, and Zhangyang Wang. Seal: Steerable
509 reasoning calibration of large language models for free. *arXiv preprint arXiv:2504.07986*, 2025.

510

511 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu,
512 Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for $2+3=?$ on the overthinking
513 of o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024b.

514

515 Jeffrey Cheng and Benjamin Van Durme. Compressed chain of thought: Efficient reasoning through
516 dense representations. *arXiv preprint arXiv:2412.13171*, 2024.

517

518 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
519 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to
520 solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

521

522 Kwesi Cobbina and Tianyi Zhou. Where to show demos in your prompt: A positional bias of
523 in-context learning. *arXiv preprint arXiv:2507.22887*, 2025.

524

525 Mengru Ding, Hanmeng Liu, Zhizhang Fu, Jian Song, Wenbo Xie, and Yue Zhang. Break the chain:
526 Large language models can be shortcut reasoners. *arXiv preprint arXiv:2406.06580*, 2024.

527

528 Gongfan Fang, Xinyin Ma, and Xinchao Wang. Thinkless: Llm learns when to think. *arXiv preprint*
529 *arXiv:2505.13379*, 2025.

530

531 Yichao Fu, Junda Chen, Yonghao Zhuang, Zheyu Fu, Ion Stoica, and Hao Zhang. Reasoning without
532 self-doubt: More efficient chain-of-thought through certainty probing. In *ICLR 2025 Workshop*
533 on Foundation Models in the Wild, 2025.

534

535 Kanishk Gandhi, Denise Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and
536 Noah D Goodman. Stream of search (sos): Learning to search in language, 2024. *URL*
537 <https://arxiv.org/abs/2404.03683>, 2, 2024.

538

539 Zorik Gekhman, Eyal Ben David, Hadas Orgad, Eran Ofek, Yonatan Belinkov, Idan Szpektor,
540 Jonathan Herzig, and Roi Reichart. Inside-out: Hidden factual knowledge in llms. *arXiv preprint*
541 *arXiv:2503.15299*, 2025.

542

543 Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. Did aristotle
544 use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of*
545 *the Association for Computational Linguistics*, 9:346–361, 2021.

540 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 541 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 542 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

543

544 Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong
 545 Tian. Training large language models to reason in a continuous latent space. *arXiv preprint*
 546 *arXiv:2412.06769*, 2024.

547

548 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,
 549 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for
 550 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint*
 551 *arXiv:2402.14008*, 2024.

552

553 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 554 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv*
 555 *preprint arXiv:2103.03874*, 2021.

556

557 Shijue Huang, Hongru Wang, Wanjun Zhong, Zhaochen Su, Jiazhan Feng, Bowen Cao, and Yi R
 558 Fung. Adactrl: Towards adaptive and controllable reasoning via difficulty-aware budgeting. *arXiv*
 559 *preprint arXiv:2505.18822*, 2025a.

560

561 Yao Huang, Huanran Chen, Shouwei Ruan, Yichi Zhang, Xingxing Wei, and Yinpeng Dong.
 562 Mitigating overthinking in large reasoning models via manifold steering. *arXiv preprint*
 563 *arXiv:2505.22411*, 2025b.

564

565 Rob J Hyndman and Yanan Fan. Sample quantiles in statistical packages. *The American Statistician*,
 566 50(4):361–365, 1996.

567

568 Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec
 569 Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. Openai o1 system card. *arXiv*
 570 *preprint arXiv:2412.16720*, 2024.

571

572 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
 573 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
 574 evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.

575

576 Guochao Jiang, Guofeng Quan, Zepeng Ding, Ziqin Luo, Dixuan Wang, and Zheng Hu. Flashthink:
 577 An early exit method for efficient reasoning. *arXiv preprint arXiv:2505.13949*, 2025.

578

579 Mingyu Jin, Qinkai Yu, Dong Shu, Haiyan Zhao, Wenyue Hua, Yanda Meng, Yongfeng Zhang, and
 580 Mengnan Du. The impact of reasoning step length on large language models. *arXiv preprint*
 581 *arXiv:2401.04925*, 2024.

582

583 Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. C3ot: Generating shorter chain-of-thought
 584 without compromising effectiveness. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 585 volume 39, pp. 24312–24320, 2025.

586

587 Aayush Karan and Yilun Du. Reasoning with sampling: Your base model is smarter than you think.
 588 *arXiv preprint arXiv:2510.14901*, 2025.

589

590 Aviral Kumar, Vincent Zhuang, Rishabh Agarwal, Yi Su, John D Co-Reyes, Avi Singh, Kate Baumli,
 591 Shariq Iqbal, Colton Bishop, Rebecca Roelofs, et al. Training language models to self-correct via
 592 reinforcement learning. *arXiv preprint arXiv:2409.12917*, 2024.

593

594 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph
 595 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 596 serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*,
 597 pp. 611–626, 2023a.

598

599 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph
 600 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
 601 serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*,
 602 pp. 611–626, 2023b.

594 Ayeong Lee, Ethan Che, and Tianyi Peng. How well do llms compress their own chain-of-thought?
595 a token complexity approach. *arXiv preprint arXiv:2503.01141*, 2025.
596

597 Yaniv Leviathan, Matan Kalman, and Yossi Matias. Fast inference from transformers via speculative
598 decoding. In *International Conference on Machine Learning*, pp. 19274–19286. PMLR, 2023.

599 Peiji Li, Kai Lv, Yunfan Shao, Yichuan Ma, Linyang Li, Xiaoqing Zheng, Xipeng Qiu, and Qipeng
600 Guo. Fastmcts: A simple sampling strategy for data synthesis. *arXiv preprint arXiv:2502.11476*,
601 2025.

602

603 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
604 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth
605 International Conference on Learning Representations*, 2023a.

606

607 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan
608 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth
609 International Conference on Learning Representations*, 2023b.

610 Weizhe Lin, Xing Li, Zhiyuan Yang, Xiaojin Fu, Hui-Ling Zhen, Yaoyuan Wang, Xianzhi Yu,
611 Wulong Liu, Xiaosong Li, and Mingxuan Yuan. Trimr: Verifier-based training-free thinking
612 compression for efficient test-time scaling. *arXiv preprint arXiv:2505.17155*, 2025a.

613

614 Zhengkai Lin, Zhihang Fu, Ze Chen, Chao Chen, Liang Xie, Wenxiao Wang, Deng Cai, Zheng
615 Wang, and Jieping Ye. Controlling thinking speed in reasoning models. *arXiv preprint
616 arXiv:2507.03704*, 2025b.

617

618 Chengzhi Liu, Zhongxing Xu, Qingyue Wei, Juncheng Wu, James Zou, Xin Eric Wang, Yuyin Zhou,
619 and Sheng Liu. More thinking, less seeing? assessing amplified hallucination in multimodal
620 reasoning models. *arXiv preprint arXiv:2505.21523*, 2025a.

621

622 Ruikang Liu, Yuxuan Sun, Manyi Zhang, Haoli Bai, Xianzhi Yu, Tiezheng Yu, Chun Yuan, and
623 Lu Hou. Quantization hurts reasoning? an empirical study on quantized reasoning models. *arXiv
624 preprint arXiv:2504.04823*, 2025b.

625

626 Sheng Liu, Haotian Ye, Lei Xing, and James Zou. In-context vectors: Making in context learning
627 more effective and controllable through latent space steering. *arXiv preprint arXiv:2311.06668*,
628 2023.

629

630 Sheng Liu, Haotian Ye, Lei Xing, and James Zou. Reducing hallucinations in vision-language
631 models via latent space steering. *arXiv preprint arXiv:2410.15778*, 2024.

632

633 Chenwei Lou, Zewei Sun, Xinnian Liang, Meng Qu, Wei Shen, Wenqi Wang, Yuntao Li, Qing-
634 ping Yang, and Shuangzhi Wu. Adacot: Pareto-optimal adaptive chain-of-thought triggering via
635 reinforcement learning. *arXiv preprint arXiv:2505.11896*, 2025.

636

637 Yijia Luo, Yulin Song, Xingyao Zhang, Jiaheng Liu, Weixun Wang, GengRu Chen, Wenbo Su, and
638 Bo Zheng. Deconstructing long chain-of-thought: A structured reasoning optimization framework
639 for long cot distillation. *arXiv preprint arXiv:2503.16385*, 2025.

640

641 Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning
642 models can be effective without thinking. *arXiv preprint arXiv:2504.09858*, 2025a.

643

644 Wenjie Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. Reasoning
645 models can be effective without thinking. *arXiv preprint arXiv:2504.09858*, 2025b.

646

647 Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke
648 Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. s1: Simple test-time
649 scaling. *arXiv preprint arXiv:2501.19393*, 2025.

650

651 Sania Nayab, Giulio Rossolini, Marco Simoni, Andrea Saracino, Giorgio Buttazzo, Nicolamaria
652 Manes, and Fabrizio Giacomelli. Concise thoughts: Impact of output length on llm reasoning and
653 cost. *arXiv preprint arXiv:2407.19825*, 2024.

648 OpenAI. GPT-3.5 Turbo Models. <https://platform.openai.com/docs/models/gpt-3-5>, 2023. Accessed: 2025-11-22.
649
650

651 OpenCompass. Aime 2025, February 2025. URL <https://huggingface.co/datasets/opencompass/AIME2025>.
652

653 Samuel J Paech. Eq-bench creative writing benchmark v3. <https://github.com/EQ-bench/creative-writing-bench>, 2025.
654
655

656 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Di-
657 rani, Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a bench-
658 mark. In *First Conference on Language Modeling*, 2024.
659

660 Matthew Renze and Erhan Guven. The benefits of a concise chain of thought on problem-solving in
661 large language models. In *2024 2nd International Conference on Foundation and Large Language
662 Models (FLLM)*, pp. 476–483. IEEE, 2024.
663

664 Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai
665 Wang, Zhaoxiang Liu, and Shiguo Lian. Dast: Difficulty-adaptive slow-thinking for large reason-
666 ing models. *arXiv preprint arXiv:2503.04472*, 2025.
667

668 Leheng Sheng, An Zhang, Zijian Wu, Weixiang Zhao, Changshuo Shen, Yi Zhang, Xiang Wang,
669 and Tat-Seng Chua. On reasoning strength planning in large reasoning models. *arXiv preprint
670 arXiv:2506.08390*, 2025.
671

672 Oscar Skean, Md Rifat Arefin, Dan Zhao, Niket Patel, Jalal Naghiyev, Yann LeCun, and Ravid
673 Shwartz-Ziv. Layer by layer: Uncovering hidden representations in language models. *arXiv
674 preprint arXiv:2502.02013*, 2025.
675

676 Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute
677 optimally can be more effective than scaling parameters for reasoning. In *The Thirteenth Inter-
678 national Conference on Learning Representations*, 2025.
679

680 Gaurav Srivastava, Shuxiang Cao, and Xuan Wang. Towards reasoning ability of small language
681 models. *arXiv preprint arXiv:2502.11569*, 2025.
682

683 DiJia Su, Hanlin Zhu, Yingchen Xu, Jiantao Jiao, Yuandong Tian, and Qingqing Zheng. Token
684 assorted: Mixing latent and text tokens for improved language model reasoning. *arXiv preprint
685 arXiv:2502.03275*, 2025a.
686

687 Jinyan Su and Claire Cardie. Thinking fast and right: Balancing accuracy and reasoning length with
688 adaptive rewards. *arXiv preprint arXiv:2505.18298*, 2025.
689

690 Jinyan Su, Jennifer Healey, Preslav Nakov, and Claire Cardie. Between underthinking and over-
691 thinking: An empirical study of reasoning length and correctness in llms. *arXiv preprint
692 arXiv:2505.00127*, 2025b.
693

694 Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu,
695 Andrew Wen, Shaochen Zhong, Hanjie Chen, et al. Stop overthinking: A survey on efficient
696 reasoning for large language models. *arXiv preprint arXiv:2503.16419*, 2025.
697

698 Hanshi Sun, Momin Haider, Ruiqi Zhang, Huitao Yang, Jiahao Qiu, Ming Yin, Mengdi Wang, Peter
699 Bartlett, and Andrea Zanette. Fast best-of-n decoding via speculative rejection. *Advances in
700 Neural Information Processing Systems*, 37:32630–32652, 2024.
701

702 Zhongxiang Sun, Qipeng Wang, Haoyu Wang, Xiao Zhang, and Jun Xu. Detection and miti-
703 gation of hallucination in large reasoning models: A mechanistic perspective. *arXiv preprint
704 arXiv:2505.12886*, 2025.
705

706 Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025. URL
707 <https://qwenlm.github.io/blog/qwq-32b/>.
708

709 Chenlong Wang, Yuanning Feng, Dongping Chen, Zhaoyang Chu, Ranjay Krishna, and Tianyi
710 Zhou. Wait, we don't need to "wait"! removing thinking tokens improves reasoning efficiency.
711 *arXiv preprint arXiv:2506.08343*, 2025a.
712

702 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 703 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 704 *arXiv preprint arXiv:2203.11171*, 2022.

705 Yiming Wang, Pei Zhang, Siyuan Huang, Baosong Yang, Zhuosheng Zhang, Fei Huang, and Rui
 706 Wang. Sampling-efficient test-time scaling: Self-estimating the best-of-n sampling in early de-
 707 coding. *arXiv preprint arXiv:2503.01422*, 2025b.

708 Yue Wang, Qiuwei Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu,
 709 Juntao Li, Zhuosheng Zhang, et al. Thoughts are all over the place: On the underthinking of
 710 o1-like llms. *arXiv preprint arXiv:2501.18585*, 2025c.

711 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 712 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in
 713 neural information processing systems*, 35:24824–24837, 2022.

714 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
 715 Pierrick Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. Huggingface’s transformers:
 716 State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019.

717 Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. Tokenskip: Controllable
 718 chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*, 2025.

719 Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. Chain of draft: Thinking faster by writing
 720 less. *arXiv preprint arXiv:2502.18600*, 2025a.

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

723 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,
 724 Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint
 725 arXiv:2505.09388*, 2025a.

726 Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Qiaowei Li, Zheng Lin, Li Cao,
 727 and Weiping Wang. Dynamic early exit in reasoning models. *arXiv preprint arXiv:2504.15895*,
 728 2025b.

729 Junjie Yang, Ke Lin, and Xing Yu. Think when you need: Self-adaptive chain-of-thought learning.
 730 *arXiv preprint arXiv:2504.03234*, 2025c.

731 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik
 732 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Ad-
 733 vances in neural information processing systems*, 36:11809–11822, 2023.

734 Ping Yu, Jing Xu, Jason Weston, and Ilia Kulikov. Distilling system 2 into system 1. *arXiv preprint
 735 arXiv:2407.06023*, 2024.

736 Hang Yuan, Bin Yu, Haotian Li, Shijun Yang, Christina Dan Wang, Zhou Yu, Xueyin Xu, Weizhen
 737 Qi, and Kai Chen. Not all tokens are what you need in thinking. *arXiv preprint arXiv:2505.17827*,
 738 2025.

739 Weizhe Yuan, Ilia Kulikov, Ping Yu, Kyunghyun Cho, Sainbayar Sukhbaatar, Jason Weston, and Jing
 740 Xu. Following length constraints in instructions, 2024. URL <https://arxiv.org/abs/2406.17744>,
 741 2024.

742 Linan Yue, Yichao Du, Yizhi Wang, Weibo Gao, Fangzhou Yao, Li Wang, Ye Liu, Ziyu Xu, Qi Liu,
 743 Shimin Di, et al. Don’t overthink it: A survey of efficient r1-style large reasoning models. *arXiv
 744 preprint arXiv:2508.02120*, 2025.

745 Jinghan Zhang, Xiting Wang, Fengran Mo, Yeyang Zhou, Wanfu Gao, and Kunpeng Liu. Entropy-
 746 based exploration conduction for multi-step reasoning. *arXiv preprint arXiv:2503.15848*, 2025a.

747 Nan Zhang, Yusen Zhang, Prasenjit Mitra, and Rui Zhang. When reasoning meets compression:
 748 Benchmarking compressed large reasoning models on complex reasoning tasks. *arXiv preprint
 749 arXiv:2504.02010*, 2025b.

756 Shengjia Zhang, Junjie Wu, Jiawei Chen, Changwang Zhang, Xingyu Lou, Wangchunshu Zhou,
757 Sheng Zhou, Can Wang, and Jun Wang. Othink-r1: Intrinsic fast/slow thinking mode switching
758 for over-reasoning mitigation. *arXiv preprint arXiv:2506.02397*, 2025c.

759 Shimao Zhang, Yu Bao, and Shujian Huang. Edt: Improving large language models' generation by
760 entropy-based dynamic temperature sampling. *arXiv preprint arXiv:2403.14541*, 2024.

762 Zhen Zhang, Xuehai He, Weixiang Yan, Ao Shen, Chenyang Zhao, Shuohang Wang, Yelong Shen,
763 and Xin Eric Wang. Soft thinking: Unlocking the reasoning potential of llms in continuous
764 concept space. *arXiv preprint arXiv:2505.15778*, 2025d.

765 Yuqi Zhu, Jia Li, Ge Li, YunFei Zhao, Zhi Jin, and Hong Mei. Hot or cold? adaptive temperature
766 sampling for code generation with large language models. In *Proceedings of the AAAI Conference
767 on Artificial Intelligence*, volume 38, pp. 437–445, 2024.

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 CONTENTS
811

812	A Supplementary Motivation and Evidence	17
813	A.1 From Overthinking to Underthinking	17
814	A.2 Confidence-Length Association	17
815	A.3 Markov Persistence of Confidence States	18
816	A.4 Distributional Heterogeneity in Model Confidence	20
817	A.5 Encoding of Confidence Signals in Latent Representations	20
818	A.6 Confidence as Evidence under Vocabulary Coverage Gaps	21
819		
820	B Method Details	24
821	B.1 Explicit Modeling of Overthinking and Underthinking	24
822	B.2 Confidence-Based Steering Vector Extraction	26
823	B.3 Model Behavior-Based Dynamic Control Function	27
824		
825	C Additional Experimental Results and Ablations	29
826	C.1 Ablation on Individual Axes	29
827	C.2 Ablation on Gating Mechanism	30
828	C.3 Cross-Domain and Cross-Difficulty Transferability	30
829	C.4 Pass@k and Avg@k Performance Analysis	31
830	C.5 Performance Variation under Different Confidence Distributions	33
831	C.6 Semantic Change and Creativity Analysis	34
832	C.7 Confidence Characteristics of Overthinking and Underthinking	36
833	C.8 Performance comparison with TrimR and Flashthink	36
834	C.9 Balanced Thinking with Dynamic Temperature	37
835	C.10 Additional Prototype Construction Strategies	38
836		
837	D Details on Experimental Settings	39
838		
839	E Details on Benchmarks	40
840		
841	F Details on Prompts	41
842		
843	G Details on Prompt-Based Approaches	42
844		
845	H Detailed Discussion of Related Works	43
846		
847	I Efficiency Analysis	45
848		
849	J The Use of Large Language Models	46
850		
851	K Ethics Statement	46
852		
853	L Reproducibility Statement	46
854		
855	M Case Study	47




Figure 6: Mitigating Overthinking Without Inducing Underthinking. Evaluation on the *MATH* dataset with DeepSeek-R1-Distill-Qwen-1.5B. Each panel plots output length vs. reference length (words), restricted to reference ≤ 300 and output ≤ 8000 . Top: correct examples. Bottom: incorrect examples. Competing methods reduce overthinking at the cost of underthinking, whereas **Our Method** mitigates overthinking without inflating underthinking.

A SUPPLEMENTARY MOTIVATION AND EVIDENCE

A.1 FROM OVERTHINKING TO UNDERTHINKING

Why do many anti-overthinking techniques backfire as underthinking? Empirically, the chain of thought (CoT) length and model performance are positively correlated, so aggressively truncating or penalizing long chains can excise necessary reasoning and degrade accuracy (Jin et al., 2024). Token-complexity theory (Lee et al., 2025) further posits an intrinsic minimum token budget for success. Enforcing uniformly short budgets or early termination pushes more instances below this threshold, yielding concise but wrong outputs. Moreover, the reasoning-boundary framework (Chen et al., 2024a) shows that optimal CoT length and reasoning path selection are task-dependent, thus global length controls disregard this heterogeneity and may steer reasoning trajectories outside the feasible region for a given task.

As shown in Fig. 6, on the *MATH* dataset with DEEPSEEK-R1-DISTILL-QWEN-1.5B, prior methods indeed curb *overthinking* but often induce *underthinking*, manifesting as a collapse of error distributions toward short outputs. In contrast, **Our Method** adaptively identifies and modulates the reasoning process, selectively shortening reasoning chains when appropriate while preserving longer explorations necessary for challenging instances. Consequently, our approach mitigates overthinking without inducing underthinking, as evidenced by error distributions that avoid collapsing into shorter outputs.

A.2 CONFIDENCE-LENGTH ASSOCIATION

Although confidence is commonly used as a proxy for a model’s certainty, its connection to actual reasoning behavior remains ambiguous. To effectively utilize confidence for identifying suboptimal reasoning patterns or enabling adaptive control, it is essential to first establish a clear and quantifiable relationship between confidence and specific aspects of model behavior. Building upon the quantitative patterns presented in Fig. 2(b), which reveal distinct confidence signatures associated with overthinking and underthinking, we focus here on another key dimension of efficient reasoning: reasoning length.

Specifically, we measure how response length relates to step-level confidence on **MATH-500**, using four models (DeepSeek-R1-Distill-Qwen-1.5B, DeepSeek-R1-Distill-Qwen-7B, Qwen3-14B, QwQ-32B; each with $n = 500$ samples). Here, the length of a response is defined as the total number of words in the reasoning text before the first `</think>`, and this count is aligned with the list of confidence values for each step. For every answer, we compute two key quantities: (i) the *minimum* step-level confidence, and (ii) the *variance* of step-level confidence within the answer.

Since confidence values produced by non-greedy decoding tend to be heavily skewed toward the upper end of the interval $[0, 1]$ rather than following a balanced and symmetric bell-shaped distribu-

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
tion, the usual Pearson correlation is not suitable. Therefore, we use the nonparametric Spearman rank correlation to capture the relationship more reliably.

Our correlation analysis results are shown in Tab. 4. Across all four models, word count shows a clear negative association with the minimum step-level confidence (Spearman ρ values roughly between -0.62 and -0.80), and a clear positive association with the variance of step-level confidence within each answer (Spearman ρ values roughly between 0.46 and 0.70). For all reported entries, the 95% bootstrap confidence intervals do not include zero. These results provide quantitative evidence that confidence trajectories encode meaningful signals about reasoning effort and stability, supporting their use as reliable indicators for dynamic reasoning control.

Model	Words vs Min confidence			Words vs Confidence variance		
	ρ	CI_lo	CI_hi	ρ	CI_lo	CI_hi
DeepSeek-R1-Distill-Qwen-1.5B	-0.733	-0.777	-0.684	0.602	0.542	0.655
DeepSeek-R1-Distill-Qwen-7B	-0.624	-0.681	-0.560	0.458	0.385	0.527
QwQ-32B	-0.801	-0.835	-0.765	0.698	0.643	0.746
Qwen3-14B	-0.681	-0.730	-0.628	0.623	0.561	0.682

Table 4: Spearman correlations (ρ) between word count length and confidence statistics on **MATH-500**. The 95% bootstrap percentile confidence intervals (CI_lo and CI_hi, with $B = 2000$) are reported in separate columns. All correlation coefficients are significant at $p < 0.001$.

A.3 MARKOV PERSISTENCE OF CONFIDENCE STATES

Effective dynamic control over a model’s reasoning trajectory often relies on the ability to recognize its current reasoning state. Intuitively, this requires examining a contextual window of recently generated reasoning steps, i.e., a sliding window over the chain-of-thought trace. However, for complex problems such as those in AIME (AI-MO, 2024a), reasoning traces can span thousands of tokens. While a larger window might seem necessary to capture sufficient context, it introduces significant computational overhead and may obscure fine-grained shifts in reasoning behavior.

In this section, we demonstrate that the model’s confidence trajectory exhibits strong first-order Markov persistence. This finding reveals a key insight: the current reasoning state can be accurately inferred from just the immediately preceding step. Consequently, a minimal window of size two is sufficient and often preferable for capturing the essential dynamics of the reasoning process.

To formalize this, for each answer, we split the model outputs by double newlines ($\backslash n \backslash n$) and align the resulting segments with the sentence-level confidences; we only consider adjacencies that occur within complete, untruncated reasoning trajectories that contain the final answer. We then collect all adjacent confidence pairs (c_{t-1}, c_t) from the model’s reasoning traces.

We convert each confidence value into either a high state or a low state using the model-wise median threshold τ :

$$s_t = \mathbb{I}(c_t \geq \tau), \quad s_t \in \{H, L\}, \quad \tau = \text{median}\{c_t \text{ over all sentences of the model}\}.$$

Here H (high) represents $c_t \geq \tau$ and L (low) represents $c_t < \tau$. If a confidence value equals the threshold, that sentence is placed in the high state.

For each model, we form a two-by-two transition count matrix

$$\mathbf{N} = \begin{bmatrix} HH & HL \\ LH & LL \end{bmatrix},$$

where HH is the number of transitions from high to high and HL is the number of transitions from high to low.

By normalizing each row, we obtain the corresponding transition probabilities:

$$P(H \rightarrow H) = \frac{HH}{HH + HL}, \quad P(H \rightarrow L) = \frac{HL}{HH + HL},$$

$$P(L \rightarrow H) = \frac{LH}{LH + LL}, \quad P(L \rightarrow L) = \frac{LL}{LH + LL}.$$

972 To measure how strongly a state tends to be followed by the same state, we use the odds ratio
 973
 974

$$975 \quad \text{OR} = \frac{HH \cdot LL}{976 \cdot HL \cdot LH}.$$

$$977$$

$$978$$

979 We evaluate statistical significance using a two-sided Fisher exact test applied to \mathbf{N} . When all cells
 980 of the matrix are positive, we also report the approximate ninety-five percent confidence interval for
 981 the odds ratio (Woolf method):
 982

$$983 \quad \log(\widehat{\text{OR}}) \pm 1.96 \sqrt{\frac{1}{HH} + \frac{1}{HL} + \frac{1}{LH} + \frac{1}{LL}}, \quad \text{CI}_{95\%} = \exp(\cdot).$$

$$984$$

$$985$$

$$986$$

$$987$$

988 For easier interpretation, we additionally provide the same state rate
 989

$$990 \quad \text{SameRate} = \frac{HH + LL}{991 \cdot HH + HL + LH + LL}.$$

$$992$$

$$993$$

$$994$$

995 All values reported in our results, including transition probabilities, odds ratios, confidence intervals,
 996 and significance levels, are computed using this median-based thresholding procedure.
 997

998 From Tab. 5, all four models show clear evidence of strong like-to-like persistence in confidence
 999 when using the model-wise median threshold. The transition probabilities for remaining in the
 1000 same state, $P(H \rightarrow H)$ and $P(L \rightarrow L)$, are both larger than the probabilities of switching to the
 1001 opposite state. The overall same state rate satisfies $\text{SameRate} > 0.5$, the Fisher exact tests give
 1002 $p < 0.001$, and the diagonal odds ratios are consistently greater than one with ninety-five percent
 1003 confidence intervals that do not include one. Taken together, these results provide strong support
 1004 for the presence of first-order Markov persistence, which reflects a clear tendency for the confidence
 1005 state to remain stable from one step to the next.
 1006

Model	$P(H \rightarrow H)$	$P(L \rightarrow L)$	SameRate	OR	CI.lo	CI.hi	Sig.
DeepSeek-R1-Distill-Qwen-1.5B	0.666	0.665	0.666	3.96	3.83	4.10	***
DeepSeek-R1-Distill-Qwen-7B	0.653	0.651	0.652	3.51	3.38	3.65	***
QwQ-32B	0.670	0.673	0.672	4.19	4.03	4.35	***
Qwen3-14B	0.657	0.659	0.658	3.70	3.56	3.86	***

1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014 **Table 5:** Adjacent state persistence in step-level confidence using a *median* threshold for binarization. Rows
 1015 report the same state transition probabilities and diagonal odds ratios (OR) with ninety-five percent Woolf
 1016 confidence intervals. Significance codes (Sig.): * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two sided Fisher
 1017 exact test).
 1018
 1019

1020 This observation directly guides the design of our sliding window. Based on this insight, we set
 1021 the window size to $w = 2$ instead of a larger value. A window of length two records each pair of
 1022 adjacent states and therefore captures the transition patterns $P(H \rightarrow L)$ and $P(L \rightarrow H)$ without any
 1023 loss of information. This choice keeps detection lag to a minimum and prevents brief reversals from
 1024 being averaged away. When the model begins to drift away from its current reasoning direction,
 1025 for example, when it moves into an overthinking regime, the adjacent transition window allows the
 intervention strength to increase immediately.

1026
1027

A.4 DISTRIBUTIONAL HETEROGENEITY IN MODEL CONFIDENCE

1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041

In this study, our goal is to design a dynamic control function that leverages stepwise confidence and its variance during the reasoning process to enable real-time, adaptive regulation of model behavior. We aim to make this control mechanism highly adaptable, i.e., capable of fully realizing our proposed concept of balanced thinking and delivering accurate, smooth, and responsive control. However, greater adaptability inherently involves introducing additional parameters, which raises concerns about the generalization ability of such a control function across diverse models. To address this, we analyze and compare confidence distributions across multiple models to investigate whether a universal hyperparameter configuration can be identified.

1042
1043
1044
1045
1046
1047
1048
1049

As shown in Fig. 7, we visualize confidence distributions across reasoning steps for multiple models on the same dataset. The three models based on the QWEN2 family, namely QwQ-32B, DEEPSEEK-R1-DISTILL-QWEN-7B, and DEEPSEEK-R1-DISTILL-QWEN-1.5B, display broadly similar distributional patterns, although each model still exhibits its own characteristic shape. In contrast, the QWEN3-14B model, built upon the QWEN3 family, exhibits a clearly different confidence profile compared with the QWEN2 family. These observations are consistent with previous findings that the LLAMA-3.1-NEMOTRON-NANO-8B model tends to operate in a uniformly low confidence regime throughout its reasoning process (Yang et al., 2025b).

1050
1051
1052
1053
1054
1055
1056

These distributional differences highlight the difficulty of designing a single set of hyperparameters capable of effectively generalizing across various models. Consequently, this motivates our proposed *model behavior-based dynamic control function fitting* approach in Sec. 3.4, which automatically derives parameters tailored to the unique confidence behaviors of each model. By leveraging this behavior-aware strategy, our method eliminates the need for manual hyperparameter tuning, thereby ensuring robust adaptability and broad applicability across diverse model architectures and confidence profiles.

1057
1058

A.5 ENCODING OF CONFIDENCE SIGNALS IN LATENT REPRESENTATIONS

1059
1060
1061
1062
1063
1064
1065

To enable dynamic control over a model’s reasoning behavior through confidence-aware steering, it is crucial to understand how confidence manifests in the model’s internal representations. Since hidden states directly encode the evolving behavioral dynamics of a transformer during reasoning, they provide a natural substrate for both analyzing and manipulating the model’s certainty. In this section, we demonstrate that confidence is not merely correlated with, but systematically and predominantly linearly encoded in hidden states. This insight underpins an automated approach for identifying the most effective layers to target in steering interventions.

1066
1067
1068
1069

Confidence is discernible in hidden layers. From Eq. 1, we obtain stepwise confidence values. We also extract the hidden state $\mathbf{H}_s^{(i)}$ of the token following the delimiter `\n\n`. This gives rise to a direct mapping between hidden state and confidence:

$$\mathbf{H}_s^{(i)} \longmapsto c_s.$$

1070
1071
1072
1073
1074
1075
1076
1077

Thus, the layer- i hidden state for sentence s corresponds to its confidence c_s . As shown in Fig. 8, we apply t-SNE to project the hidden states $\mathbf{H}_s^{(i)}$ into a two-dimensional space and color each point according to its corresponding confidence c_s . Clear clusters emerge: high-confidence and low-confidence representations form visibly distinct regions, with this separation becoming even more pronounced in the mid-late layer embeddings. These observations indicate that confidence acts as a discernible signal in the hidden space, providing direct empirical support for our subsequent linear probing analysis.

1078
1079

Linear probing of the confidence signal. Building on the above observations, we employ a linear probing approach to examine how confidence is encoded in the hidden layers and to assess the




Figure 7: KDE (Confidence) for QwQ-32B, QwQ-32B, QwQ-32B, and DeepSeek-R1-Distill-Qwen-7B and 1.5B on MATH-500.

strength of this linear relationship. Formally, the mapping is expressed as:

$$c_s = \mathbf{w}^{(i)} \mathbf{H}_s^{(i)} + b^{(i)}.$$

Here, $\mathbf{w}^{(i)}$ and $b^{(i)}$ denote the parameters of the linear model.

Because hidden layer representations typically have several thousand dimensions, using a linear head directly may lead to overfitting and unstable estimation. To address this, we apply a standard dimensionality reduction method, PCA, to project the hidden states into a low-dimensional subspace, and then study the relationship between the reduced representations and the corresponding confidence values. The detailed statistics before and after PCA are summarized in Tab. 6.

$$c_s = \mathbf{w}^{(i)} \mathbf{Z}_s^{(i)} + b^{(i)}, \quad \mathbf{H}_s^{(i)} \longmapsto \mathbf{Z}_s^{(i)}.$$

Here, $\mathbf{Z}_s^{(i)}$ denotes the low-dimensional representation obtained from the original hidden state $\mathbf{H}_s^{(i)}$.

The overall probe analysis pipeline is illustrated in Fig. 10. We employ a standard ridge regression approach to estimate the parameters $\mathbf{w}^{(i)}$ and $b^{(i)}$.

$$\min_{\mathbf{w}^{(i)}, b^{(i)}} \sum_s \left(c_s - \mathbf{w}^{(i)} \mathbf{Z}_s^{(i)} - b^{(i)} \right)^2 + \lambda \left\| \mathbf{w}^{(i)} \right\|_2^2.$$

After fitting the ridge-based linear probe, we obtain the predicted confidence values as follows.

$$\hat{c}_s = \mathbf{w}^{(i)} \mathbf{Z}_s^{(i)} + b^{(i)}.$$

We assess how accurately confidence can be predicted from the hidden representations by computing the coefficient of determination R^2 . The formulation is given below. A higher R^2 value, approaching 1, indicates that confidence is more readily linearly decodable from the representations of the corresponding layer.

$$R^2 = 1 - \frac{\sum_s (c_s - \hat{c}_s)^2}{\sum_s (c_s - \bar{c})^2}.$$

Automated steering layer selection using R^2 . We evaluate the relationship between the confidence values c_s and the hidden state $\mathbf{H}_s^{(i)}$ across all layers of each model, as shown in Fig. 9. A consistent pattern emerges: the coefficient of determination R^2 typically reaches its maximum in the middle to late layers, indicating that c_s is more easily linearly decodable from $\mathbf{H}_s^{(i)}$ in these layers. Since steering fundamentally operates through linear shifts in the hidden space, we select the layer with the highest R^2 as the steering layer. The entire procedure is fully automated, allowing the system to identify the optimal steering layer without manual intervention. In principle, choosing this layer minimizes the additional noise introduced by the steering operation.

Model	Original Dim	PCA Dim	Retained Var.
DeepSeek-R1-Distill-Qwen-1.5B	1 536	64	0.889 3
DeepSeek-R1-Distill-Qwen-7B	3 584	64	0.857 3
QwQ-32B	5 120	64	0.833 5
Qwen3-14B	5 120	64	0.900 0

Table 6: Cumulative explained variance retained by PCA ($k = 64$) across models. Higher retained variance indicates a stronger low-dimensional linear structure amenable to probing.

A.6 CONFIDENCE AS EVIDENCE UNDER VOCABULARY COVERAGE GAPS

A growing number of efficient reasoning methods mitigate overthinking by relying on predefined keyword vocabularies. Representative strategies include *targeted suppression* of specific tokens (e.g., NOWAIT (WANG ET AL., 2025A)), *latent-space guidance* that steers the model away from




Figure 8: t-SNE projections of hidden states sampled immediately after `\n\n`; colors encode sentence-level confidence (0–1). (a) DeepSeek-R1-Distill-Qwen-1.5B — layers 5, 20, 28. (b) DeepSeek-R1-Distill-Qwen-7B — layers 5, 20, 28. (c) Qwen3-14B — layers 5, 20, 35. (d) QwQ-32B — layers 5, 20, 58.

states prone to emitting certain tokens (e.g., SEAL (Chen et al., 2025), MANIFOLD STEERING (Huang et al., 2025b)), and *cue-word-driven early stopping*, which treats designated trigger terms as checkpoints to terminate the reasoning process when appropriate (e.g., FLASHTHINKING (JIANG ET AL., 2025), TRIMR (LIN ET AL., 2025A), DEER (YANG ET AL., 2025B),

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197




Figure 9: Layer-wise linear probe decodability (R^2) of confidence. Mid-to-late layers achieve the highest scores, motivating the steering-layer choice.

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241




Figure 10: Linear-confidence probe with PCA. Given an input sequence, we freeze the language model (LM) and extract layer-wise hidden states. The representations are projected by PCA and passed to a linear predictor to estimate step-level confidence. We repeat this procedure across layers to obtain layer-wise test-set R^2 scores.

DYNASOR-CoT (FU ET AL., 2025)). However, the fundamental reason why such lexical interventions improve reasoning behavior remains poorly understood. In this section, we aim to uncover the mechanism behind their effectiveness by aggregating keyword vocabularies from representative methods and analyzing their relationship with model confidence. Our key finding is that vocabulary-based strategies are, in essence, incomplete approximations of confidence-based control: they capture only the most frequent lexical manifestations of low-confidence reasoning, while missing a broader spectrum of uncertainty signals.

Category	Vocabulary
NoWait (suppress)	wait, alternatively, hmm, but, however, alternative, another, check, double-check, oh, maybe, verify, other, again, now, ah, any
SEAL—Transition	alternatively, think differently, another way, another approach, another method, another solution, another strategy, another technique
SEAL—Reflection	wait, verify, make sure, hold on, think again, 's correct, 's incorrect, let me check, seems right

Table 7: Unified vocabularies for NoWait (keyword suppression) and SEAL transition/reflection cues.

To investigate this, we first compile the keyword sets used by NOWAIT and SEAL as illustrative examples (see Tab. 7). Notably, these lexical items function as surface markers of the model’s epistemic uncertainty. Using DEEPSPEEK-R1-DISTILL-QWEN-7B on MATH-500, we compute a sentence-level confidence at each reasoning step and project it to all words appearing in that

1242 sentence. As summarized in Tab. 8, the vast majority of these words are associated with confi-
 1243 dences below the model’s overall mean ($\bar{c} = 0.8162$). This pattern suggests a straightforward
 1244 interpretation: vocabulary-based interventions such as NOWAIT and SEAL primarily suppress *low-
 1245 confidence modes* of the model’s reasoning, rather than targeting particular semantics per se.
 1246

1247 As illustrated in Fig. 11, both SEAL and NoWait reliably elevate the model’s confidence along
 1248 the reasoning trajectory. This empirical pattern corroborates our analysis: NOWAIT achieves the
 1249 effect by suppressing the emission of high-frequency lexical markers associated with low confi-
 1250 dence, whereas SEAL steers the hidden representations away from states that tend to produce low-
 1251 confidence, high-frequency sentences. In essence, both methods act by attenuating the model’s
 1252 low-confidence modes.

1253 However, vocabulary-driven heuristics do not, by
 1254 themselves, capture the model’s *low-confidence
 1255 modes*. In practice, such methods identify only a sub-
 1256 set of *high-frequency lexical correlates* of low confi-
 1257 dence, leaving a substantial long-tail of equally in-
 1258 formative cues outside the predefined lists and thus
 1259 unmeasured. As illustrated in Tab. 8, we enum-
 1260 erate several representative omissions that most exist-
 1261 ing approaches fail to account for. Consequently,
 1262 confidence-based approaches systematically surface
 1263 the model’s low-confidence modes—irrespective of
 1264 their lexical realization.

1264 This comprehensive extraction provides a principled basis for subsequent research to diagnose and
 1265 mitigate overthinking, enabling more complete coverage than vocabulary-driven heuristics. Look-
 1266 ing forward, a fruitful research agenda is to pursue a confidence-based line of work. One direction
 1267 is to treat low-confidence states as actionable checkpoints for early exit, developing calibrated cri-
 1268 teria and adaptive halting policies to further improve the accuracy–efficiency trade-off of early-exit
 1269 models. Another is to analyze the relationship between semantic (meaning-level) uncertainty and
 1270 model-internal confidence estimates, thereby deepening our understanding of—and ultimately miti-
 1271 gating—both overthinking and underthinking behaviors.
 1272

1273 B METHOD DETAILS


1274 In this section, we provide a more detailed introduction to the technical details of REBALANCE
 1275 presented in Sec. 3. First, in Sec. B.1, we formally define these reasoning modes and propose an
 1276 explicit, confidence-based modeling paradigm to quantitatively distinguish between redundant rea-
 1277 soning and premature conclusion. Building upon this foundation, Sec. B.2 presents the *Confidence-
 1278 Based Steering Vector Extraction*, where we leverage hidden-state representations to derive di-
 1279 rectional steering vectors that guide LRM’s reasoning trajectory towards optimal decision-making
 1280 boundaries. Finally, Sec. B.3 details our *Model Behavior-Based Dynamic Control Function*, which
 1281 dynamically modulates steering strength according to real-time confidence and variance metrics. By
 1282 integrating these three sequential components, our framework achieves a robust and adaptive control
 1283 mechanism, effectively balancing exploration and commitment in the reasoning processes of LRM.
 1284

1285 B.1 EXPLICIT MODELING OF OVERTHINKING AND UNDERTHINKING

1286 **Formal Definition.** Let the reasoning trajectory inside `<think>...</think>` be split into steps
 1287 $S_1, \dots, S_{s_{\max}}$ by the double newline delimiter `\n\n` introduced in Sec. 2.1. Let $r_{\leq s}$ denote the
 1288 partial reasoning up to step s . If the model is forced to stop after step s and produce a conclusion
 1289 from $r_{\leq s}$, it induces a distribution over answers which we denote by π_s . Let $d_s = \arg \max \pi_s$ be
 1290 the predicted conclusion under a fixed decoding rule. Define the stability index
 1291

$$1293 s^* = \min \{ s : d_{s'} = d_s \text{ for all } s' \geq s \text{ and } d_s \text{ is correct} \}.$$

1294 A trajectory exhibits *overthinking* if it continues generating steps after s^* . Conversely, A trajectory
 1295 exhibits *underthinking* if it stops at step s with an incorrect d_s while there exists $s' > s$ such that $d_{s'}$



1296 **Figure 11:** KDE (Confidence) - Origin vs
 1297 NoWait vs SEAL.

1296
 1297
 1298
 1299
 1300
 1301
 1302
 1303
 1304
 1305
 1306
 1307
 1308
 1309
 1310
 1311
 1312
 1313
 1314
 1315
 1316
 1317
 1318
 1319
 1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

Table 8: Low-confidence lexicon (items from the curated vocabulary are marked with * and colored red).

Block A				Block B				Block C				Block D			
Word/Phrase	Confidence	Count	Word/Phrase	Confidence	Count	Word/Phrase	Confidence	Count	Word/Phrase	Confidence	Count	Word/Phrase	Confidence	Count	Word/Phrase
think differently *	0.6067	49	alternative *	0.6375	13	think again *	0.6537	115	another approach *	0.6549	209				
sorry	0.5990	1	confusing	0.6620	71	confused	0.6690	101	maybe *	0.6740	3,012				
another method *	0.6781	78	another way *	0.6821	417	alternatively *	0.6844	1,857	hold on *	0.6881	491				
however *	0.6914	68	i think	0.6920	866	in summary	0.6940	10	i suspect	0.6950	1				
verify (variants) *	0.6971	354	verify *	0.6988	339	any *	0.7115	729	i'm not sure	0.7040	60				
let me think	0.7050	647	i guess	0.7120	43	it seems	0.7120	101	probably	0.7120	77				
but *	0.7092	6,386	double-check *	0.7161	339	not sure	0.7180	154	Let me check *	0.7223	384				
check *	0.7228	1,093	other *	0.7290	776	again *	0.7347	734	wait *	0.7325	2,223				
anyway	0.7270	6	another solution *	0.7357	10	possibly	0.7490	2	hmm *	0.7710	1,728				
i believe	0.7700	11	's correct *	0.7821	288	wow	0.7820	13	's incorrect *	0.7821	288				
it looks like	0.7830	17	seems right *	0.7919	92	phew	0.7900	1	now *	0.8182	1,734				
oh *	0.8187	40	mandatory	0.5948	5	easiest	0.5948	6	perspective	0.5977	8				
summarizing	0.4997	14	rely	0.5066	10	overlooked	0.5093	6	systematically	0.5342	5				
uses	0.5359	5	partially	0.5389	11	layer	0.5447	5	reason	0.5461	5				
uniquely	0.5468	7	summed	0.5470	7	separate	0.6460	75	annual	0.5602	5				
trick	0.5612	5	misinterpret	0.5649	5	generating	0.5702	9	crucial	0.5759	6				
substitutions	0.5763	8	consist	0.5797	7	systems	0.5809	10	neat	0.5809	5				
despite	0.5811	9	careful	0.5817	21	shake	0.5818	10	redo	0.5819	8				
clear	0.5859	29	accept	0.5866	15	discriminants	0.5869	6	haven	0.5871	18				
quickly	0.5895	5	diametrically	0.5899	5	misapplied	0.5914	8	homogeneous	0.5937	6				
extends	0.5979	5	consuming	0.6028	7	obvious	0.6033	7	altered	0.6038	6				
worried	0.6046	7	periodicity	0.6055	11	theory	0.6067	10	verification	0.6072	11				
interpreting	0.6073	11	absolutely	0.6076	12	translation	0.6082	5	schedule	0.6098	5				
elsewhere	0.6100	9	designed	0.6123	8	shapes	0.6148	10	may	0.6150	15				
interpolation	0.6159	7	hard	0.6162	8	necessary	0.6162	42	unfolding	0.6164	27				
concerning	0.4954	2	uncertain	0.5537	1	tangled	0.6371	6	mentally	0.6378	10				
surprised	0.6059	1	relief	0.6126	1	allows	0.6379	9	mix	0.6380	31				
doubt	0.6559	1	unsure	0.6699	1	solidify	0.6394	5	arrive	0.6401	10				
concerned	0.7315	3	surprising	0.7318	5	precisely	0.6407	13	effectively	0.6413	30				
nervous	0.8295	1	verifying	0.6416	13	mixing	0.6416	11	offset	0.6416	12				
triplets	0.6420	7	complicate	0.6422	38	extensions	0.6426	10	handshake	0.6427	8				
version	0.6430	8	eighteen	0.6432	16	abstract	0.6433	9	paper	0.6433	5				
nature	0.6435	7	clarify	0.6442	44	shaking	0.6442	7	overcomplicating	0.6442	45				
exploit	0.6443	5	cell	0.6444	5	safe	0.6444	13	interpreted	0.6445	9				
pyramid	0.6446	6	assumptions	0.6448	6	crossing	0.6451	7				

would be correct. These definitions capture redundant reasoning beyond the earliest stable correct decision and premature commitment before sufficient exploration.

Explicit Modeling with Confidence. We instantiate the above definition using the sequence of stepwise confidence $\{c_s\}$ and confidence variance $\{v_s\}$ as defined in Sec. 2.1, where $v_s = \text{Var}(c_s; \mathcal{W}_s)$ and \mathcal{W}_f is a sliding window. We can determine two-sided quantile thresholds from a small-scale seen dataset. Let $Q_c(q)$ and $Q_v(q)$ be the empirical q -quantile of $\{c_s\}$ and $\{v_s\}$. Choose $0 < q_L < q_H < 1$ and define

$$\tau_c^L = Q_c(q_L), \quad \tau_c^H = Q_c(q_H), \quad \tau_v^L = Q_v(q_L), \quad \tau_v^H = Q_v(q_H).$$

A step is tagged as *low-confidence* if $c_s \leq \tau_c^L$ and *high-confidence* if $c_s \geq \tau_c^H$. Similarly, A step is tagged as *high-variance* if $v_s \geq \tau_v^H$ and *low-variance* if $v_s \leq \tau_v^L$. We then define the sets

$$\mathcal{O} = \{s : c_s \leq \tau_c^L \text{ and } v_s \geq \tau_v^H\}, \quad \mathcal{U} = \{s : c_s \geq \tau_c^H \text{ and } v_s \leq \tau_v^L\}.$$

As observed in Fig. 2(b), high variance reflects frequent switching across reasoning paths and often co-occurs with low confidence; thus, we treat \mathcal{O} as a proxy for overthinking. Persistently high confidence with low variance indicates stable yet potentially premature commitment, making \mathcal{U} a proxy for underthinking. Steps that fall outside both sets are considered to reflect a normal state and are excluded from subsequent analyses.

B.2 CONFIDENCE-BASED STEERING VECTOR EXTRACTION

Building upon the explicit modeling paradigm, we propose deriving steering vectors from deep-layer hidden representations to guide LRMAs away from undesirable reasoning modes. These vectors are efficiently obtained via a one-pass collection performed only once per model prior to deployment, eliminating additional computation during actual use.

One-Pass Data Collection. We prepare a small-scale seen dataset $\mathcal{D}_{\text{seed}}$ and run the model once per prompt. When the model generates a delimiter `\n\n`, the next token marks the first token of a new step. At this token, we save deep-layer hidden states $\mathbf{h}_{\ell, t_s^{(1)}}$ for step index s and selected layers ℓ , chosen via a probing method maximizing confidence separability on a single dataset but shared across all datasets (see Appendix A.5). The first token of a step serves as a compact representation of the step mode for two reasons. First, it typically encodes the intent that sets the direction of the step (e.g., *wait* or *alternatively*) (Yang et al., 2025b), and due to the causal mask, all later tokens in the step condition on it. Second, deep layers show stronger distinguishability between the two reasoning modes in our empirical study, consistent with Gekhman et al. (2025); Skean et al. (2025).

Steering Vector Extraction Using the tags from sets \mathcal{O} and \mathcal{U} , we form latent distributions for the overthinking and underthinking modes. For each selected layer ℓ , we obtain mode prototypes by averaging the hidden states obtained from the one-pass data collection

$$\boldsymbol{\mu}_\ell^{\mathcal{O}} = \frac{1}{|\mathcal{O}|} \sum_{s \in \mathcal{O}} \mathbf{h}_{\ell, t_s^{(1)}}, \quad \boldsymbol{\mu}_\ell^{\mathcal{U}} = \frac{1}{|\mathcal{U}|} \sum_{s \in \mathcal{U}} \mathbf{h}_{\ell, t_s^{(1)}}.$$

The difference between the two prototypes defines a steering vector for the ℓ -th layer

$$\mathbf{v}_\ell = \frac{\boldsymbol{\mu}_\ell^{\mathcal{O}} - \boldsymbol{\mu}_\ell^{\mathcal{U}}}{\|\boldsymbol{\mu}_\ell^{\mathcal{O}} - \boldsymbol{\mu}_\ell^{\mathcal{U}}\|_2}.$$

This vector encodes the transition direction in latent space from underthinking toward overthinking, with its negation representing the reverse.

During inference, we inject the steering vector solely at each step’s first token. Specifically, we modify the deep hidden state by

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

$$\tilde{\mathbf{h}}_{\ell, t_s^{(1)}} = \mathbf{h}_{\ell, t_s^{(1)}} + \alpha_\ell \mathbf{v}_\ell.$$

Let $t_s^{(1)}$ be its position and let $\alpha_\ell \in \mathbb{R}$ be a scalar steering weight that controls strength λ_ℓ and direction λ_ℓ .

$$\alpha_\ell = \lambda_\ell \delta_\ell, \quad \lambda_\ell \geq 0, \quad \delta_\ell \in \{+1, -1\}.$$

Setting $\delta_\ell = +1$ pushes the state away from underthinking and increases exploration. Setting $\delta_\ell = -1$ pushes the state away from overthinking and facilitates the model to coverage to a reasonable reasoning path. The values of λ_ℓ and δ_ℓ will be determined by the dynamic control function introduced in the later section. Conceptually, the two prototypes act as boundaries of the model’s reasoning process. Our goal is to keep the stepwise state between these boundaries so that the model reasons efficiently with balanced thinking.

B.3 MODEL BEHAVIOR-BASED DYNAMIC CONTROL FUNCTION

Inputs differ in difficulty, and the model’s reasoning state evolves over time. To keep the trajectory between the overthinking and underthinking boundaries, we set the steering weight α online as a continuous function of the current state. The function takes the stepwise confidence c_s and the confidence variance v_s as inputs, and outputs a steering weight α that determines both direction δ and magnitude λ . The weight pushes the state away from the closer boundary and grows as the state approaches that boundary.

From a confidence curve to a control surface. To derive this three-dimensional surface, we first construct a simplified two-dimensional curve $f(c)$ based solely on confidence c . From the previous analysis, the steering weight α needs to transition smoothly from a minimum negative value (away from overthinking) to a maximum positive value (away from underthinking) as confidence c increases. Many functions satisfy this requirement. Here, we adopt the widely used sigmoid as an illustration.

For ease of spatial transformation, we express the sigmoid function in terms of the hyperbolic tangent:

$$\sigma(c) = \frac{1}{1 + e^{-c}} = \frac{1}{2} + \frac{1}{2} \tanh\left(\frac{c}{2}\right).$$

Our goal is to spatially transform this sigmoid to align precisely with the overthinking and underthinking boundaries. After transformation, the function becomes:

$$f(c) = a + b \tanh(k(c + m)),$$

where a , b , k , and m represent spatial transformation parameters, which can be obtained by fitting.

However, as detailed in Appendix A.4, confidence distributions vary significantly across models, making it difficult to find universally applicable parameters. Thus, we propose a *model behavior-based* fitting method. This method adaptively determines these parameters based on model-specific behavior, using the previously collected stepwise confidence c_s and confidence variance v_s from the one-pass data collection without additional computational cost.

Specifically, after the one-pass collection, we obtain hidden-state distributions and corresponding prototypes for overthinking and underthinking, from which we derive a steering vector \mathbf{v} . Since the steering vector connects prototypes that represent their respective hidden-state distributions, the steering strength can be interpreted as the displacement of these distributions along the vector direction. Therefore, adjusting the magnitude of this displacement enables us to capture specific behavioral characteristics of the model, allowing tailored data point generation.

To illustrate this, consider first the alleviation of overthinking. Suppose the hidden-state distribution of overthinking is bounded. The aggressive displacement is defined as the minimal distance required

1458 to shift all points within this distribution beyond its boundary. A more moderate displacement, how-
 1459 ever, moves only the overthinking prototype outside the boundary, defining the moderate distance.
 1460 We explicitly anchor this through specific behavioral criteria:

1461

- 1462 • Anchor A1: At $c = \tau_c^L$, the steering should yield a negative moderate displacement, effec-
 1463 tively guiding the state away from overthinking.
- 1464 • Anchor A2: At $c = \tau_c^H$, the steering is set to zero, as the state is considered to lie within
 1465 the normal confidence region, and thus no additional steering is applied.

1466 In contrast, mitigating underthinking poses unique challenges. Experimental observations indicate
 1467 that LRM_s may also consistently exhibit high confidence during normal reasoning, making direct
 1468 numerical quantification of overconfidence, a defining characteristic of underthinking tendency, dif-
 1469 ficult. Therefore, we adopt a conservative mitigation approach. Recognizing that greater disper-
 1470 sion in confidence distributions corresponds to larger distances between prototypes measured by the
 1471 norm of the steering vector, we select aggressive and moderate displacements for underthinking as
 1472 small proportional fractions of this norm. This proportional strategy reduces underthinking with-
 1473 out adversely affecting normal reasoning and ensures scalability across diverse LRM_s’ confidence
 1474 distributions. This approach is anchored by:

1475

- 1476 • Anchor A3: At $c = 1$, the maximum normalized confidence, the steering provides a posi-
 1477 tive moderate displacement to mitigate excessive confidence and counteract underthinking.

1478 We obtain these moderate targets from model behavior in latent space. Let prototype distance $d^{\text{prot}} =$
 1479 $\|\mu^O - \mu^U\|_2$ and let $s_s = \mathbf{v}^\top \mathbf{h}_{t_s^{(1)}}$. Define a separating threshold along \mathbf{v} by $t = \frac{1}{2} \mathbf{v}^\top (\mu^O + \mu^U)$.
 1480 The *moderate* distance for overthinking is

$$1483 \quad d^{O,m} = \mathbf{v}^\top \mu^O - t,$$

1485 which moves the overthinking prototype to the boundary. The *aggressive* distance for overthinking is
 1486

$$1488 \quad d^{O,a} = \max_{s \in \mathcal{O}} (s_s) - t,$$

1490 which moves all overthinking steps past the boundary. For underthinking, we adopt a conservative
 1491 rule since normal reasoning can also show sustained high confidence. We set

$$1493 \quad d^{U,m} = \rho_U^m d^{\text{prot}}, \quad d^{U,a} = \rho_U^a d^{\text{prot}},$$

1495 with constants $0 < \rho_U^m < \rho_U^a$. These distances scale with the separation between prototypes and
 1496 adapt across models. Because \mathbf{v} has unit norm, a displacement of size d along \mathbf{v} corresponds to a
 1497 steering magnitude d in the hidden space. We therefore fit f so that

$$1499 \quad f(\tau_c^L) = -d^{O,m}, \quad f(\tau_c^H) = 0, \quad f(1) = +d^{U,m}.$$

1501 We solve for b and k by least squares on these anchors. This yields a smooth curve that produces
 1502 negative weights at low confidence and positive weights at high confidence, with zero at the center.

1503

1504 **Lifting to a variance aware surface.** We now incorporate variance to obtain a two-input control
 1505 $g(c, v)$. The idea is to keep the sign and basic shape from $f(c)$ while increasing the magnitude near
 1506 the two high-risk regions. The overthinking region is $c \leq \tau_c^L$ with $v \geq \tau_v^H$. The underthinking
 1507 region is $c \geq \tau_c^H$ with $v \leq \tau_v^L$. However, abrupt steering strength changes at region boundaries
 1508 are undesirable. To mitigate this, we define smooth gates that approach one inside each region and
 1509 decay to zero outside

$$1511 \quad \psi_O(c, v) = \sigma\left(\frac{\tau_c^L - c}{\eta_c}\right) \sigma\left(\frac{v - \tau_v^H}{\eta_v}\right), \quad \psi_U(c, v) = \sigma\left(\frac{c - \tau_c^H}{\eta_c}\right) \sigma\left(\frac{\tau_v^L - v}{\eta_v}\right),$$




Figure 12: Left: heatmap of $g(c_s, v_s)$ with v on a log scale; Right: 3D surface of $g(c_s, v_s)$. Dashed lines mark q_{25}^c, q_{75}^c and q_{25}^v, q_{75}^v .

where $\sigma(x) = 1/(1 + e^{-x})$. The parameters $\eta_c > 0$ and $\eta_v > 0$ control the width of the transition and smooth the change near region boundaries. We then modulate the amplitude of f by replacing the moderate distances with aggressive distances inside the corresponding region. Let

$$B(c, v) = B^m + (B^{O,a} - B^m) \psi_O(c, v) + (B^{U,a} - B^m) \psi_U(c, v),$$

with B^m chosen by the fit of f , $B^{O,a} = d^{O,a}$, and $B^{U,a} = d^{U,a}$. The final control surface is

$$g(c, v) = \text{sign}(c - \tau_c^H) B(c, v) \tanh(k |c - \tau_c^H|).$$

For fixed v , the map is monotone in c . For fixed c , the magnitude increases smoothly as v enters the overthinking or underthinking region. The sign follows the confidence side so that the weight always pushes away from the nearer boundary.

Online steering. At step s we set

$$\alpha_s = g(c_s, v_s), \quad \lambda_s = |\alpha_s|, \quad \delta_s = \text{sign}(\alpha_s).$$

These values plug into the injection rule and separate direction and magnitude as defined earlier. The procedure is training-free and uses only statistics that are already computed online. It adapts across models through the behavior-based distances and across inputs through the gates on (c_s, v_s) . The result is a continuous controller that keeps the trajectory between the two mode boundaries and allocates more steering when the state drifts toward either boundary.

Function Visualization. We visualize the fitted mapping $g(c_s, v_s)$ for DEEPSEEK-R1-DISTILL-QWEN-1.5B. Darker regions—characterized by *high variance* and *low confidence*—indicate where the *overthinking penalty* is strongest. Conversely, the lighter, positive region at *low variance* and *high confidence* marks where the *underthinking penalty* is strongest.

C ADDITIONAL EXPERIMENTAL RESULTS AND ABLATIONS

C.1 ABLATION ON INDIVIDUAL AXES

We study univariate effects via axis-wise ablations of $g(c_s, v_s)$: define $g_c(c_s) := g(c_s, \bar{v})$ and $g_v(v_s) := g(\bar{c}, v_s)$, where \bar{v} and \bar{c} denote the means of confidence variance v and stepwise confidence c estimated on the extraction set. As shown in the first three rows of Tab. 9, both univariate

1566 1567 1568 1569 1570 1571	Method	Math500		GSM8K		Olympiad	
		Pass@1 \uparrow	#Tokens \downarrow	Pass@1 \uparrow	#Tokens \downarrow	Pass@1 \uparrow	#Tokens \downarrow
		Full [$g(c, v)$]	83.0	3474	78.3	765	43.9
	$g(c)$	82.6 \downarrow	3387 \downarrow	78.1 \downarrow	890 \uparrow	41.7 \downarrow	6612 \downarrow
	$g(v)$	76.6 \downarrow	3596 \uparrow	78.0 \downarrow	658 \downarrow	39.6 \downarrow	6876 \downarrow

1572
1573
1574
1575
1576
Table 9: Axis-wise ablations on the R1-1.5B backbone across difficulty levels. We analyze performance changes on three math benchmarks with varying difficulty: **Math500** (medium), **GSM8K** (easy), and **Olympiad** (hard). Metrics are Pass@1 accuracy (%) and generated token numbers. Arrows indicate change relative to **Full** [$g(c, v)$]: Acc. \uparrow increase, \downarrow decrease; Tokens \downarrow decrease, \uparrow increase.

1577 1578 1579	Gating Functions	MATH-500		GSM8K		Olympiad		GPQA-D	
		Pass@1 \uparrow	#Tokens \downarrow						
DeepSeek-R1-Distill-Qwen-1.5B									
1580	Baseline	79.6	4516	76.0	1018	41.2	8785	17.1	8727
1581	Sigmoid Gate	83.0	3474	78.3	765	43.9	7235	21.7	6902
1582	Linear Gate	83.6	3600	79.2	800	42.5	6794	17.7	7019
1583	Hard-Step Gate	81.2	3830	79.7	820	41.8	7528	17.2	6537
1584	Polynomial-Fitted Gate	82.8	3508	79.0	798	44.2	7059	20.7	6161
1585	ReLU-Shaped Gate	84.0	3277	80.1	818	43.0	6875	19.7	6966
DeepSeek-R1-Distill-Qwen-7B									
1586	Baseline	89.8	3699	89.2	1098	56.1	7590	33.8	7392
1587	Sigmoid Gate	92.6	2903	91.6	912	57.0	6321	39.4	5180
1588	Linear Gate	91.2	3113	90.0	936	56.2	6303	43.0	5589
1589	Hard-Step Gate	92.0	3115	90.0	909	57.8	6319	40.4	5448
1590	Polynomial-Fitted Gate	91.6	2932	91.3	913	56.0	6206	38.9	5768
	ReLU-Shaped Gate	91.4	3054	91.7	928	57.0	6362	39.9	5758

1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
Table 10: Performance comparison of different gating functions across benchmarks. We evaluate both **DeepSeek-R1-Distill-Qwen-1.5B** and **DeepSeek-R1-Distill-Qwen-7B** on four reasoning datasets. Metrics are Pass@1 accuracy (%) and the number of generated tokens. Different gating mechanisms exhibit notable variations in both accuracy and efficiency.

variants degrade REBALANCE performance, indicating that the bivariate form $g(c, v)$ provides finer-grained and more effective control than single-variable schemes.

C.2 ABLATION ON GATING MECHANISM

Our control surface design is originally built upon a smoothly parameterized sigmoid gating function, which is a widely adopted choice in prior work due to its simplicity and smoothness properties. The primary role of this gating mechanism, as outlined in Sec. 3.4 and further detailed in Appendix B.3, is to produce a smoother fitted surface and avoid abrupt transitions in steering strength during inference.

To investigate the sensitivity of our method to the specific form of the gating function, we conduct an ablation study by replacing the default sigmoid gate with several alternative gating strategies. As shown in Tab. 10, we evaluate four variants: a linear gate, a hard-step gate, a polynomial-fitted gate, and a ReLU-shaped gate. Results across multiple models, difficulty levels, and domains indicate that all variants achieve effective reasoning performance, with only minor differences observed across datasets. This demonstrates that our approach is robust to the particular choice of gating function.

Notably, while the sigmoid function yields strong empirical performance relative to baseline, it was selected primarily for its common usage rather than through extensive engineering optimization. Our experiments suggest that with additional tuning or more sophisticated gate designs, performance could potentially be further improved. However, since the precise design of the gating function is not the central focus of this work, we retain the standard sigmoid as the default for simplicity and reproducibility.

C.3 CROSS-DOMAIN AND CROSS-DIFFICULTY TRANSFERABILITY

1620 **Cross-domain transferability.** To validate the cross-domain transferability of ReBalance, we
 1621 conduct experiments on DeepSeek-R1-Distill-Qwen-1.5B. Specifically, we extract steering vectors
 1622 and confidence statistics from coding (LiveCodeBench) and commonsense (StrategyQA) tasks and
 1623 transfer them to the mathematics domain (MATH-500). We also report the quartiles of the confi-
 1624 dence distribution (q_{25c} , q_{75c}) and variance distribution (q_{25v} , q_{75v}) derived from each extraction
 1625 dataset. For readability, confidence quartiles are scaled by 100 and variance quartiles by 1000.

1626 As observed in Tab. 11, although the confidence statistics differ to some extent across domains,
 1627 ReBalance achieves strong performance in all cases. Notably, extractions from StrategyQA can
 1628 even achieve higher Pass@1 scores compared to those from MATH within the same domain.

1630
 1631 **Cross-difficulty transferability.** We categorize commonly used mathematical reasoning datasets
 1632 by difficulty level: easy (GSM8K, AMC23), medium (MATH/MATH-500), and hard (AIME24,
 1633 AIME25, Olympiad). Using the medium-difficulty datasets as an intermediary, we examine two
 1634 transfer directions, easy-to-medium and hard-to-medium, to explore how changes in task difficulty
 1635 affect the transfer performance of ReBalance. The experimental results on DeepSeek-R1-Distill-
 1636 Qwen-1.5B are shown in Tab. 11.

1637 However, data distributions vary across different datasets, which may introduce confounding factors
 1638 beyond just difficulty. Moreover, it is challenging to precisely quantify the difficulty gaps between
 1639 datasets, posing significant challenges to the analysis. Therefore, we leverage the ground-truth
 1640 difficulty grading within the MATH dataset and conduct an analysis based on QwQ-32B (Tab. 11),
 1641 performing an in-distribution, fine-grained difficulty classification from Level 1 to Level 5 within
 1642 the same dataset.

1643 Based on the experimental results, we have the following observations. Firstly, the higher the diffi-
 1644 culty of the extraction dataset, the lower the confidence and the higher the variance. This reflects the
 1645 model’s broader and more frequent exploration of reasoning paths, aligning with the decision rule
 1646 we proposed in Eq. 5, where confidence serves as an indicator. Secondly, Extraction datasets of easy
 1647 difficulty prioritize token reduction, while those of hard difficulty prioritize accuracy improvement.
 1648 We believe this occurs because easy extraction datasets exhibit higher confidence and lower vari-
 1649 ance; thus, according to Eq. 5, the range classified as overthinking is wider, and the underthinking
 1650 range is narrower. This results in more frequent triggering of the overthinking criterion, causing the
 1651 dynamic control function to preferentially suppress redundant reasoning. The opposite occurs for
 1652 harder datasets.

1653 Therefore, we suggest using extraction datasets of medium difficulty (e.g., MATH) in practical ap-
 1654 plications to achieve an optimal accuracy-efficiency trade-off.

1656 C.4 PASS@K AND AVG@K PERFORMANCE ANALYSIS

1657
 1658 The core issue addressed by ReBalance is balanced thinking, i.e., mitigating overthinking while
 1659 simultaneously preventing underthinking. Here, “underthinking” refers to cases where the model
 1660 is inherently capable of solving a problem but produces an incorrect answer due to insufficient
 1661 reasoning. According to the definition of Pass@ k , this metric effectively grants the model multiple
 1662 reasoning attempts and expanded exploration space, counting a question as correct if any one of the
 1663 k sampled solutions is successful. Therefore, theoretically, for problems that meet the underthinking
 1664 criterion, which are those that the model is truly capable of solving, the likelihood of obtaining
 1665 a correct answer approaches certainty as the number of samples increases, eventually converging
 1666 toward the model’s capability ceiling (Karan & Du, 2025). At this point, the impact of underthinking
 1667 on accuracy becomes negligible, and ReBalance primarily serves to alleviate overthinking.

1668 To validate the above analysis, we evaluate Pass@20 over 20 samples and measure the average token
 1669 length of generated sequences on the few-example datasets AMC23, AIME24, and AIME25. The
 1670 results are shown in Tab. 12.

1671 As can be seen, consistent with our analysis, ReBalance significantly reduces reasoning length with-
 1672 out any degradation in model accuracy. This demonstrates that our proposed confidence indicator
 1673 accurately characterizes overthinking and underthinking. Thanks to this precise characterization,
 ReBalance prunes only genuinely redundant reasoning steps rather than essential or effective ones,

Dataset	Confidence Distribution				MATH-500	
	q25c	q75c	q25v	q75v	Pass@1↑	#Tokens↓
QwQ-32B						
MATH	70.6	92.0	0.4	7.3	95.2	3662
Level-1	74.8	92.8	0.3	5.2	94.8	3447
Level-2	72.4	92.7	0.3	6.9	94.8	3623
Level-3	71.0	92.1	0.4	6.9	95.2	3678
Level-4	69.6	92.0	0.4	7.4	95.0	3696
Level-5	69.2	91.1	0.4	7.7	95.4	3720
DeepSeek-R1-Distill-Qwen-1.5B						
MATH	66.8	93.9	0.5	11.0	83.0	3474
GSM8K	76.9	94.4	0.3	6.7	80.6	3221
AMC23	61.8	94.2	0.4	10.4	82.8	3277
AIME24	63.2	92.1	0.5	9.8	83.2	3568
AIME25	58.9	90.2	0.4	9.3	84.0	3796
Olympiad	58.1	84.2	0.5	9.5	83.6	3862
GPQA	56.5	82.8	0.5	9.5	81.4	4260
LiveCodeBench	62.4	91.2	0.5	6.2	82.0	3482
StrategyQA	61.2	85.8	0.4	8.3	83.4	3667

Table 11: Performance Variation under Difficulty-Conditioned Control Surfaces. We investigate performance shifts induced by control surfaces derived from datasets of varying difficulty. For each difficulty tier, we extract the steering vectors and fit the associated control surface, and subsequently evaluate them on DeepSeek-R1-Distill-Qwen-1.5B and QwQ-32B. The results indicate systematic differences in Rebalance behavior, suggesting that dataset difficulty plays a non-trivial role in shaping the resulting control dynamics.

Method	AMC23		AIME2024		AIME2025	
	Pass@20↑	#Tokens↓	Pass@20↑	#Tokens↓	Pass@20↑	#Tokens↓
DeepSeek-R1-Distill-Qwen-1.5B						
Baseline	97.5	7430	63.3	10645	43.3	10447
ReBalance	97.5	4832	63.3	9243	46.7	8652
DeepSeek-R1-Distill-Qwen-7B						
Baseline	100.0	6021	76.7	11249	66.7	11379
ReBalance	100.0	5264	76.7	8563	66.7	9019
Qwen3-14B						
Baseline	100.0	7331	90.0	11367	80.0	12717
ReBalance	100.0	5124	90.0	9247	80.0	10381
QwQ-32B						
Baseline	100.0	7034	86.7	14121	83.3	13386
ReBalance	100.0	5651	86.7	9853	83.3	11586

Table 12: Pass@20 Performance and Token Efficiency. ReBalance preserves accuracy while consistently reducing the token usage of baseline models.

enabling lossless and efficient sequence compression even when the model operates near its capability ceiling.

Following prior work (Jaech et al., 2024; Guo et al., 2025), we report Avg@4 on the large-scale MATH-500 dataset. For the few-sample benchmarks AMC23, AIME24, and AIME25, we also include Avg@16 together with Avg@4 to evaluate the model’s average case reasoning performance under both low-sampling settings and medium-sampling settings. As shown in Tab. 13 and 14, we have the following observations:

- **Inter-dataset variability in randomness.** As reflected by the standard deviations of Avg@k and Tok@k, datasets with very few samples exhibit substantially higher vari-

Model	Avg@1	Tok@1	Avg@4 ± Std	Tok@4 ± Std
DeepSeek-R1-Distill-Qwen-1.5B	79.6	4516	79.6 ± 0.004	4620 ± 100.9
w/ ReBalance	83.0	3474	83.3 ± 0.007	3553 ± 55.5
DeepSeek-R1-Distill-Qwen-7B	89.8	3699	89.4 ± 0.007	3703 ± 56.4
w/ ReBalance	92.6	2903	92.6 ± 0.001	2894 ± 38.1
Qwen3-14B	93.8	4470	93.8 ± 0.001	4550 ± 38.8
w/ ReBalance	94.0	3641	94.1 ± 0.001	3674 ± 44.6

Table 13: Multi-sample evaluation on MATH-500.

Model	Avg@1	Tok@1	Avg@4 ± Std	Tok@4 ± Std	Avg@16 ± Std	Tok@16 ± Std
AIME24						
DeepSeek-R1-Distill-Qwen-1.5B						
Baseline	23.3	12596	21.7 ± 0.02	12897 ± 830.7	19.6 ± 0.03	12931 ± 830.7
ReBalance	36.7	9040	34.2 ± 0.04	9179 ± 718.4	35.6 ± 0.05	9179 ± 602.6
DeepSeek-R1-Distill-Qwen-7B						
Baseline	40.0	13994	41.7 ± 0.03	13636 ± 434.2	41.3 ± 0.04	13764 ± 418.2
ReBalance	56.7	9012	55.0 ± 0.02	8664 ± 896.5	57.9 ± 0.05	9167 ± 620.9
Qwen3-14B						
Baseline	66.7	10888	70.0 ± 0.02	11488 ± 188.1	67.1 ± 0.06	11174 ± 307.9
ReBalance	73.3	9464	71.7 ± 0.02	9627 ± 115.3	73.1 ± 0.02	9613 ± 112.1
AIME25						
DeepSeek-R1-Distill-Qwen-1.5B						
Baseline	16.7	14556	15.0 ± 0.02	14107 ± 354.3	15.4 ± 0.02	14589 ± 416.3
ReBalance	30.0	8140	27.5 ± 0.03	8447 ± 447.0	27.5 ± 0.03	8723 ± 626.9
DeepSeek-R1-Distill-Qwen-7B						
Baseline	26.7	13778	28.3 ± 0.04	12340 ± 308.3	27.3 ± 0.04	12192 ± 308.3
ReBalance	40.0	9227	40.0 ± 0.02	8813 ± 609.3	40.0 ± 0.02	9241 ± 463.2
Qwen3-14B						
Baseline	56.7	13125	55.0 ± 0.05	12900 ± 241.7	54.4 ± 0.06	12457 ± 298.7
ReBalance	56.7	11057	57.5 ± 0.06	11023 ± 213.1	57.3 ± 0.07	11013 ± 224.3
AMC23						
DeepSeek-R1-Distill-Qwen-1.5B						
Baseline	55.0	8990	53.8 ± 0.01	8616 ± 302.6	54.1 ± 0.02	8637 ± 538.8
ReBalance	80.0	5216	80.6 ± 0.04	4730 ± 255.3	80.0 ± 0.04	4729 ± 433.5
DeepSeek-R1-Distill-Qwen-7B						
Baseline	75.0	6898	74.9 ± 0.02	6297 ± 520.2	74.8 ± 0.02	6297 ± 335.1
ReBalance	95.0	4767	93.8 ± 0.02	4115 ± 423.4	92.2 ± 0.04	4226 ± 394.4
Qwen3-14B						
Baseline	95.0	7240	91.3 ± 0.03	7244 ± 115.3	93.1 ± 0.03	6985 ± 210.3
ReBalance	100.0	5230	98.8 ± 0.02	5120 ± 98.4	97.7 ± 0.02	4848 ± 190.2

Table 14: Multi-sample evaluation on AMC23, AIME24, and AIME25.

ance—up to an order of magnitude larger than that of large-sample datasets such as MATH-500. This affects both accuracy and generated sequence length.

- **Performance stability on large-sample datasets.** On datasets with sufficient sample size, multi-sample evaluation has minimal impact on the *relative* performance between ReBalance and the baseline. The performance gap remains largely consistent across different sampling counts.
- **Performance stability on few-sample datasets.** Although both accuracy and sequence length fluctuate more significantly with increased sampling on few-sample datasets, the relative improvement of ReBalance over the baseline remains stable. Notably, under 16-sample evaluation, ReBalance yields even larger performance gains compared to the original single-sample setting across multiple datasets and models, further demonstrating its effectiveness and robustness.

C.5 PERFORMANCE VARIATION UNDER DIFFERENT CONFIDENCE DISTRIBUTIONS

To examine the generalizability of Rebalance across distinct confidence regimes within the same model, we keep both the control surface and evaluation samples fixed while varying the model’s decoding temperature. This setup allows us to assess whether Rebalance consistently yields per-

Temperature	Confidence Distribution				Baseline		Rebalance	
	q_{25}^c	q_{75}^c	q_{25}^v	q_{75}^v	Pass@1 (%)	Token	Pass@1 (%)	Token
0.2	1	1	0	0	68.4	6248	77.0	4335
0.4	0.9602	1	0	0.0006	72.6	5694	79.6	3906
0.6	0.8507	1	0	0.0039	79.4	4584	81.6	3689
0.8	0.7156	0.9845	0.0002	0.0098	81.0	4529	84.6	3485
1.0	0.5516	0.9449	0.0006	0.0164	82.4	4660	82.4	3488

Table 15: Performance of Rebalance under different temperature settings on DeepSeek-R1-Distill-Qwen-1.5B. Rebalance consistently provides higher accuracy and shorter reasoning traces from low- to high-temperature settings, demonstrating stable generalization under varying confidence distributions.




Figure 13: (a)–(d) show the mean confidence distributions of the Baseline and Rebalance models at temperatures 0.2, 0.4, 0.8, and 1.0, respectively. The trend reveals a clear temperature-dependent effect: at lower temperatures, Rebalance systematically reduces the model’s overconfident predictions, whereas at higher temperatures, it shifts the distribution upward, counteracting the underconfidence introduced by increased sampling randomness.

formance gains under different confidence distributions. As shown in Tab. 15, Rebalance behaves adaptively across decoding temperatures. At low temperatures, the model exhibits persistently high confidence, which frequently triggers the underthinking detector; Rebalance then encourages more diverse reasoning trajectories. Conversely, at higher temperatures, the model produces higher variance and lower confidence, activating the overthinking detector; in this regime, Rebalance effectively narrows the search space and accelerates convergence. As illustrated in Fig. 13, the results provide a detailed view of how Rebalance adjusts the model’s behavior across temperature regimes. At lower temperatures, Rebalance effectively suppresses overconfident predictions, thereby reducing underthinking and substantially improving the model’s accuracy.

C.6 SEMANTIC CHANGE AND CREATIVITY ANALYSIS

Even though Appendix A.6 examines the relationship between the keyword vocabulary and confidence, it remains necessary to analyze the semantics of the model’s generated outputs. To this end, we conduct a systematic semantic analysis of the DeepSeek-R1-Distill-Qwen-1.5B reasoning traces using the Transition and Reflection vocabularies introduced in SEAL Chen et al. (2025). As shown in Tab. 16, we quantify the semantic changes of the model under NoWait, NoThinking, and our method before and after steering. For readability, TF and TF-IDF scores are scaled by a factor of 1000. We observe that both reflection and transition patterns decrease to varying degrees across all methods. However, the first two baselines suffer from noticeable accuracy drops. In contrast, our method preserves a non-negligible amount of these reasoning patterns, which substantially contributes to maintaining, and in some cases even improving, the model’s accuracy. Retaining reflection and transition patterns appears to be a key factor for preserving reasoning correctness in large language models.

To evaluate whether applying Rebalance imposes any unintended drawbacks on creativity and the naturalness of the expressions, we further assess the models using the Creative Writing v3 benchmark Paech (2025). The experimental results are summarized in Tab. 17.

We evaluate four models using Claude-Sonnet-4.5 Anthropic (2024) as the judge model. For each model, we report three metrics: (1) Rubric Score, an aggregate quality score across multiple writing dimensions; (2) Elo Score, a relative writing-quality ranking calibrated with GPT-3.5-Turbo OpenAI (2023) and DeepSeek-R1-Distill-Qwen-1.5B as fixed anchors; and (3) Fine-grained Ability Scores, covering fifteen creative-writing competencies: coherent (logical consistency and structural clar-

1836	1837	1838	Method	Reflection			Transition			Performance Change	
				Word Count	TF	TF-IDF	Word Count	TF	TF-IDF	ΔPass@1	ΔTokens
DeepSeek-R1-Distill-Qwen-1.5B											
1840	Baseline	30.3	9.7	12.1	6.7	2.4	3.9	—	—	—	—
1841	NoThinking	3.9	1.7	4.6	0.9	0.4	1.3	-4.6	-64.9%	—	—
1842	NoWait	1.6	1.7	3.5	0.1	1.9×10^{-2}	7.7×10^{-2}	-1.6	-41.4%	—	—
1843	Rebalance	10.5	6.4	9.3	1.7	1.0	2.3	+3.4	-23.1%	—	—
DeepSeek-R1-Distill-Qwen-7B											
1844	Baseline	19.2	8.1	10.3	5.2	2.3	3.7	—	—	—	—
1845	NoThinking	1.8	1.2	3.1	0.1	9.6×10^{-2}	4.5×10^{-2}	-9.2	-77.4%	—	—
1846	NoWait	1.2	1.8	3.9	0.03	2.8×10^{-2}	1.8×10^{-2}	-3.0	-32.9%	—	—
1847	Rebalance	12.8	6.7	9.2	2.7	1.4	2.8	+2.8	-21.5%	—	—
Qwen3-14B											
1848	Baseline	24.0	9.2	10.8	8.7	3.3	4.4	—	—	—	—
1849	NoThinking	10.4	5.5	9.4	2.8	1.0	2.4	0.0	-40.5%	—	—
1850	NoWait	2.1	1.8	3.3	0.0	9.4×10^{-3}	4.1×10^{-2}	-1.0	-27.9%	—	—
1851	Rebalance	13.2	6.8	8.1	4.6	2.3	3.3	+0.2	-18.5%	—	—
QwQ-32B											
1852	Baseline	26.7	11.2	13.3	10.1	3.8	5.1	—	—	—	—
1853	NoThinking	26.0	11.5	13.7	9.5	3.8	5.1	0.0	-13.7%	—	—
1854	NoWait	2.5	2.4	4.1	0.1	2.2×10^{-2}	8.1×10^{-2}	-1.0	-27.9%	—	—
1855	Rebalance	22.1	10.8	10.1	5.4	2.5	3.8	+0.4	-19.3%	—	—

Table 16: Semantic statistics of *Transition* and *Reflection* vocabularies across different methods and models, and corresponding performance changes. For each method on each model, we report the total word count, term frequency (TF), and TF-IDF score for both vocabularies, as well as the change in accuracy and tokens relative to the Baseline of the same model.

ity), creativity (originality and non-templated expression), descriptive imagery (vivid, sensory-rich description), pacing (appropriate narrative flow), elegant prose (fluency and stylistic refinement), instruction following (accurate adherence to instructions), consistent voice & tone (stable narrative voice), strong dialogue (natural, character-appropriate dialogue), sentence flow (smooth transitions between sentences), show-don’t-tell (conveying meaning through scene and action rather than exposition), avoids amateurish prose (avoidance of clichés or novice patterns), emotional depth (nuanced emotional expression), avoids positivity bias (avoidance of forced optimism), avoids purple prose (restraint from overly ornate language), believable characters (psychological plausibility and consistent character voices).

It can be observed that after applying ReBalance, the models generally maintain and even improve their performance in creativity and the naturalness of expressions. The percentage of metrics that are higher than or equal to those of the original models is notably high: 71% for DeepSeek-R1-Distill-Qwen-1.5B, 88% for DeepSeek-R1-Distill-Qwen-7B, 100% for Qwen3-14B, and 65% for QwQ-32B. Notably, Qwen3-14B achieves significant improvements across all metrics after applying ReBalance. We posit that this improvement stems from ReBalance’s ability to continuously and gently guide the reasoning process, achieving a balance between overthinking and underthinking, thereby maintaining the model within effective reasoning boundaries. Consequently, the models exhibit measurable improvements in creative-writing performance, indicating that ReBalance accelerates convergence without compromising the model’s capacity for innovative or divergent thinking.

We additionally evaluate Qwen2.5-7B-Instruct and find that its creativity scores are substantially higher than those of the distilled DeepSeek-R1-Distill-Qwen-7B, further supporting our observations. Our analysis reveals that distillation systematically reduces linguistic diversity, a property that is closely tied to creativity. Likewise, task-specific fine-tuning may limit creative expression by reinforcing existing patterns at the expense of novel exploration.

In contrast, Rebalance induces a form of cognitive restructuring in the model’s internal reasoning process, which partially restores and enhances creative expressiveness. These findings highlight an

1890	1891	1892	1893	1894	1895	1896	1897	1898	1899	1900	1901	1902	1903	1904	1905	Score		Ability														
																Rubric	Elo	Coh	Crt	Img	Pac	Ele	Inst	Voi	Dia	Flo	SDT	Ama	Emo	Pos	Pur	Ch
DeepSeek-R1-Distill-Qwen-1.5B																																
Baseline	15.0	30.0	1.3	1.8	2.8	1.9	5.4	1.5	2.9	2.9	4.1	2.5	3.2	1.2	11.5	9.2	1.8															
Rebalance	15.2	57.5	1.5	1.9	2.7	2.2	5.6	1.4	2.8	3.0	4.3	2.5	3.3	1.2	11.2	9.4	1.7															
DeepSeek-R1-Distill-Qwen-7B																																
Baseline	22.9	308.9	4.0	4.4	4.8	4.8	6.6	3.1	5.0	4.2	6.6	4.0	5.4	2.7	9.9	9.6	3.4															
Rebalance	23.4	349.7	4.3	4.5	4.7	5.1	6.7	3.0	5.0	4.6	6.6	4.1	5.4	2.9	9.9	9.9	3.5															
Qwen3-14B																																
Baseline	42.6	1295.2	8.8	7.1	10.3	7.7	8.2	9.2	11.2	6.4	8.1	5.5	6.6	8.1	10.8	8.4	8.8															
Rebalance	50.5	1368.0	12.2	7.2	11.6	11.3	9.5	12.1	12.3	8.2	9.9	6.9	8.0	9.4	12.6	9.0	10.3															
QwQ-32B																																
Baseline	52.7	1438.2	12.6	8.1	12.0	11.7	9.8	13.6	12.6	8.5	9.8	8.1	8.3	9.8	11.4	9.6	10.6															
Rebalance	52.8	1442.7	12.7	7.9	12.3	11.8	9.9	12.6	12.2	8.6	9.8	7.9	8.0	9.8	12.2	9.5	10.7															
Qwen2.5-7B-Instruct																																
Baseline	33.5	877.6	7.4	5.2	7.2	7.6	8.5	6.1	8.0	6.2	8.7	4.8	6.9	4.9	9.5	10.6	6.3															
GPT3.5-Turbo																																
Baseline	52.8	1500.0	13.2	7.1	11.5	12.2	10.5	12.7	13.3	8.1	11.0	7.2	8.9	9.1	11.5	10.4	10.8															

Table 17: Creative-writing performance on Creative Writing v3 benchmark. Ability abbreviations: Coh = Coherent; Crt = Creativity; Img = Descriptive Imagery; Pac = Pacing; Ele = Elegant Prose; Inst = Instruction Following; Voi = Consistent Voice & Tone; Dia = Strong Dialogue; Flo = Sentence Flow; SDT = Show-Don’t-Tell; Ama = Avoids Amateurish Prose; Emo = Emotional Depth; Pos = Avoids Positivity Bias; Pur = Avoids Purple Prose; Ch = Believable Characters.

1911	1912	1913	1914	1915	1916	1917	Model			Normal (O-base)			Overthinking			Normal (U-base)			Underthinking			
							Conf.	Var.	Len.	Conf.	Var.	Len.	Conf.	Var.	Len.	Conf.	Var.	Len.	Conf.	Var.	Len.	
DeepSeek-R1-Distill-Qwen-1.5B	80.4	18.5	1357	78.6	21.3	2386	85.1	22.2	2909	89.7	16.7	1726										
DeepSeek-R1-Distill-Qwen-7B	90.0	11.0	2809	81.8	23.0	5995	82.4	22.0	3259	91.2	12.0	2752										
Qwen3-14B	92.6	8.7	5763	88.1	12.1	8819	85.3	11.3	6305	9.27	8.0	5743										
QwQ-32B	84.0	16.1	3712	75.4	22.1	6377	76.2	18.5	4573	78.9	1.76	3080										

Table 18: Comparison of Confidence, Variance, and Output Length Across Thinking Modes. For each model, we report statistics for Normal responses paired with Overthinking (Normal (O-base)) and Underthinking (Normal (U-base)), together with the corresponding Overthinking and Underthinking states.

important direction for future work: developing methods that improve reasoning stability without suppressing linguistic diversity or creative generation.

C.7 CONFIDENCE CHARACTERISTICS OF OVERTHINKING AND UNDERTHINKING

Fig. 2(b) in the main text highlights a core contribution of our work: confidence serves as a continuous and reliable indicator for explicitly modeling overthinking and underthinking. In Fig. 2(b), we use DeepSeek-R1-Distill-Qwen-1.5B on MATH-500 for visualization. To demonstrate the generality of this observation, we extend the experimental setup to include four model sizes (1.5B to 32B) across three distinct model families. Tab. 18 presents a comprehensive quantitative analysis of stepwise confidence and confidence variance across these models. For clarity, confidence values are scaled by 100 and variance by 1000.

We can observe that, across all models shown in Tab. 18, the results consistently align with those of Fig. 2(b): Firstly, overthinking samples exhibit lower confidence and higher variance, suggesting hesitant and repeated switching between reasoning paths, often leading to redundant reasoning; Secondly, underthinking samples show higher confidence and lower variance, and this persistently high confidence often causes the model to prematurely commit to an incorrect reasoning path, hindering thorough exploration. This observation further confirms that confidence can serve as a general indicator, broadly applicable across various large reasoning models.

C.8 PERFORMANCE COMPARISON WITH TRIMR AND FLASHTHINK

Since the official implementations of TrimR (Lin et al., 2025a) and FlashThink (Jiang et al., 2025) are not publicly available, we reproduce both methods for a fair comparison.

1944	QwQ-32B	MATH-500		AIME24		AIME25		GSM8K		AMC23		GPQA-D	
		Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓	Pass@1↑	#Tokens↓
TrimR													
1947	threshold=0.5	93.8	3830	56.7	8345	43.3	8827	93.7	1319	90.0	6055	63.1	6380
1948	threshold=0.75	94.8	4048	70.0	10235	53.3	11397	96.7	1410	87.5	6778	69.2	7312
1949	threshold=1	93.8	4241	66.7	11007	56.7	11726	95.9	1432	90.0	6890	63.7	7012
Flashtink													
1950	max tokens=16000	94.8	2854	56.7	8757	60.0	9613	96.3	1090	90.0	5200	64.7	5751
1951	max tokens=32000	94.6	2884	66.7	11390	63.3	11218	96.7	1098	95.0	5596	66.2	5982
ReBalance(ours)													
1953	max tokens=16000	95.2	3662	70.0	10350	63.3	11575	96.8	1289	95.0	6064	67.2	6296

Table 19: Performance comparison with TrimR and Flashtink on QwQ-32B. All experiments are run with the same sampling parameters.

Reproduction of TrimR. We adopt the reflection tokens given in the paper to split reasoning into fixed-interval sub-thoughts and feed them into a verifier model for answer detection. We use a streaming inference pipeline that monitors generation through OpenAI-compatible endpoints. Since the step size is unspecified, we set it to 100 tokens, consistent with the original Fig. 8. Both the overthinking-compression module (via answer convergence) and the underthinking-compression module (via budget monitoring) are enabled accordingly. In the original study, TrimR was executed on Ascend NPUs with the NPU-native Pangu-7B model serving as the verifier. Since we do not have access to Ascend NPUs, our experiments are instead conducted on NVIDIA RTX PRO 6000 GPUs and adopt Qwen2.5-7B-Instruct (Yang et al., 2025a) as the verifier, following the other configuration outlined in their paper.

In TrimR, R denotes the underthinking threshold defined in the original paper. Specifically, the reasoning process is terminated once it reaches $R\%$ of the maximum sequence length, with $R = 0.5$ used by default. Since TrimR is originally validated only on models of 32B parameters or larger, we conduct our comparison using QwQ-32B to ensure a fair assessment of its strengths. Despite TrimR’s use of an additional 7B auxiliary model, an extra inference stage, and post-hoc optimization via repetition truncation, ReBalance still achieves significantly better performance. Moreover, our analysis of the R parameter reveals that increasing R leads to notable accuracy gains for TrimR on challenging benchmarks such as AIME24 and AIME25. This observation further supports the hypothesis that existing approaches designed to mitigate overthinking can inadvertently induce underthinking, highlighting a critical trade-off in current reasoning frameworks.

Reproduction of FlashThink. For FlashThink, we follow the paper’s core procedure, i.e., segmenting the chain-of-thought using delimiter tokens and invoking a verifier model for early exit. Like TrimR, we employ a streaming inference pipeline that monitors generation via OpenAI-compatible endpoints. Each detected segment is forwarded to the verifier (Qwen2.5-7B-Instruct) using the original prompt template. The main results for TrimR and FlashThink are reported in Tab. 1, with more detailed comparisons provided in Tab. 19.

Overall, while FlashThink and TrimR effectively reduce token usage, they incur significant inference-time overhead. Frequent verification interrupts disrupt the decoding process, stall the pipeline, and degrade KV-cache efficiency; moreover, deploying a separate verifier increases system complexity. In contrast, ReBalance introduces minimal overhead, maintains uninterrupted decoding, and simultaneously reduces token consumption and latency, offering a more practical and self-contained solution.

C.9 BALANCED THINKING WITH DYNAMIC TEMPERATURE

In this work, we use steering as an illustrative example to practically implement our idea of balancing overthinking and underthinking. Notably, the concept itself is generalizable and applicable to a variety of approaches aiming at promoting balanced thinking. To support this argument, inspired by Zhang et al. (2024) and Zhu et al. (2024), we replace the dynamic steering component in ReBalance with a much simpler approach: adjusting the temperature parameter. Specifically, when overthinking is detected during inference, we reduce the temperature to avoid excessively divergent exploration leading to redundant reasoning. Conversely, when underthinking is detected, we increase the tem-

	Method	AMC23		AIME24		AIME25		Olympiad	
		Pass@1	#Tokens	Pass@1	#Tokens	Pass@1	#Tokens	Pass@1	#Tokens
DeepSeek-R1-Distill-Qwen-1.5B									
2001	Baseline	55	8990	23.3	12596	16.7	14556	41.2	8785
2002	Dynamic Temperature	75	7344	36.7	12054	23.3	11659	44.7	8538
2003	Dynamic Steering	80	5216	36.7	9040	30.0	8140	43.9	7253
DeepSeek-R1-Distill-Qwen-7B									
2004	Baseline	75	6898	40.0	13994	26.7	13778	56.1	7590
2005	Dynamic Temperature	82.5	5856	53.3	10593	36.7	10740	57.5	7498
2006	Dynamic Steering	95.0	4767	56.7	9012	40.0	9227	57.0	6321
QwQ-32B									
2008	Baseline	87.5	7021	66.7	14342	46.7	13350	66.7	8219
2009	Dynamic Temperature	92.5	6721	66.7	11202	53.3	12134	67.6	8160
2010	Dynamic Steering	95.0	6064	70.0	10350	63.3	11575	68.6	7422

Table 20: Performance comparison across three model sizes under different inference-time control methods.

perature to broaden reasoning. Due to constraints in time and computational resources, we conduct only preliminary experiments using a discrete, binary hyperparameter setting without further tuning. Still using confidence as the indicator, we set the temperature to 1.2 upon detecting underthinking, and reduce it to 0.7 upon detecting overthinking.

The comparative results are summarized in Tab. 20. Even with such a simple configuration, our balanced thinking approach achieves significant performance improvements and reductions in reasoning length. Therefore, we believe the performance can be further enhanced by introducing a model behavior-based dynamic function fitting method similar to ReBalance, which naturally enables adaptive continuous regulation without manual hyperparameter tuning.

C.10 ADDITIONAL PROTOTYPE CONSTRUCTION STRATEGIES

In Sec. 3.2 of the main text, we provide two definitions for overthinking and underthinking in Eq. 3 and Eq. 5, respectively, and select Eq. 5 as ReBalance’s explicit modeling approach. To address potential concerns, we present a comparative analysis and experimental validation of these two definitions.

Comparative analysis. Although Eq. 3 appears more concise and intuitive, it serves only as a theoretical definition of overthinking and underthinking. Its purpose is to provide a conceptual distinction between the two phenomena through a formalized expression. This definition operates at the trajectory level, classifying an entire reasoning trajectory as overthinking or underthinking by comparing it against an idealized stability index that acts as a decision boundary. Consequently, if Eq. 3 were directly used as the indicator for steering vector extraction, it would inevitably lead to the following issues.

- **Mismatch in operational granularity.** Our method aims to adaptively adjust the model’s behavior based on the real-time reasoning state at each step. This requires a step-level, fine-grained indicator to detect tendencies toward overthinking or underthinking. To avoid a mismatch in operational granularity, we therefore maintain the same step-level resolution during the steering vector extraction. However, Eq. 3 only supports trajectory-level classification of reasoning modes, and this granularity mismatch may lead to suboptimal performance.
- **Limited applicability.** The stability index requires access to ground truth, which introduces an additional dependency on labeled data. Although one could follow Lin et al. (2025a) using an external verifier to approximate ground truth by assessing the existence and equivalence of answers across consecutive sub-thoughts, this strategy still relies on the verifier’s capability, prompt design, and hyperparameters for determining answer equivalence. Consequently, it may incur extra engineering overhead and introduce potential performance bottlenecks.
- **Difficulty in capturing complex reasoning dynamics.** Eq. 3 defines the stability index only when all subsequent reasoning steps yield identical answers that exactly match the ground truth. However, existing studies have shown that the accuracy of reasoning mod-

Interval	1	2	3	4	5	6	7	8	9	10
Relative Position Range	[0.0, 0.1)	[0.1, 0.2)	[0.2, 0.3)	[0.3, 0.4)	[0.4, 0.5)	[0.5, 0.6)	[0.6, 0.7)	[0.7, 0.8)	[0.8, 0.9)	[0.9, 1.0]
Count of Stability Indices	27.41	16.45	10.96	9.87	6.58	3.29	2.85	2.19	3.95	16.45

Table 21: Distribution of stability indices over relative position intervals.

Methods	$\ S\ _2$	MATH-500 (Pass@1)	MATH-500 (#Tokens)	AIME24 (Pass@1)	AIME24 (#Tokens)
Baseline	—	79.6	4516	23.3	12596
Vector Extraction w/ Stability Index	11.4	81.8	4715	20.0	12563
Vector Extraction w/ Confidence	62.4	83.0	3474	36.7	9040

Table 22: Comparison of steering vector extraction variants.

els is not positively correlated with reasoning length (Chen et al., 2024b); on the contrary, longer reasoning sequences may introduce more hallucinations (Liu et al., 2025a). Therefore, while Eq. 3 offers a convenient and intuitive way to conceptually distinguish overthinking from underthinking, its rigid requirement makes it unsuitable as a practical indicator in real-world scenarios, given the inherent complexity and variability of actual reasoning dynamics.

Experimental validation. To quantitatively validate the above analysis, we perform steering vector extraction using Equation 3 as follows. First, consistent with existing methods, we randomly select 500 seen samples from the MATH training set and feed them into DeepSeek-R1-Distill-Qwen-1.5B. We collect the generated output sequences and record the model’s confidence at each reasoning step. To identify steps corresponding to overthinking and underthinking, following Lin et al. (2025a) and Jiang et al. (2025), we use Qwen2.5-7B-Instruct (Yang et al., 2025a) to determine whether each step contains the ground truth (see Appendix G for the prompt template). Once a step containing the ground truth is detected, it is designated as the stability index: all preceding steps are classified as underthinking, and all subsequent steps as overthinking. The labeling procedure produces 29,701 overthinking samples and 24,710 underthinking samples, showing that our partitioning strategy is reasonable and results in a well-balanced dataset.

To analyze the relative positional distribution of stability indices within the entire thinking sequence, we divide the range of relative positions (0–1) into 10 equal intervals (bins). We then count the number of occurrences of stability indices falling into each interval, as shown in Tab. 21. Notably, an interesting observation emerges here: the stability indices exhibit a bimodal distribution along the thinking sequence, with pronounced concentrations around Interval 1 and 10. These two Intervals correspond to the underthinking and overthinking behaviors of reasoning models, respectively, further providing strong empirical support for the necessity of balanced thinking.

The subsequent steps, including steering vector extraction, fitting of the dynamic control function, and dynamic steering during inference, are kept identical to those in ReBalance. The experimental results, along with the comparison of steering vector norms $\|S\|_2$, are shown in Tab. 22.

We observe that the experimental results obtained using the steering vector extracted via Eq. 5 significantly outperform those of Eq. 3 in both accuracy and efficiency, strongly validating our analysis above. Moreover, the difference in $\|S\|_2$ reveals that when switching to stability index-based extraction, the norm of the steering vector becomes substantially smaller. This indicates a blurring of the boundary between overthinking and underthinking, further explaining the performance gap between the two methods.

D DETAILS ON EXPERIMENTAL SETTINGS

Benchmarks. Evaluation is conducted on *mathematics reasoning* datasets: MATH-500 (Lightman et al., 2023b), AIME24 (AI-MO, 2024a), AIME25 (OpenCompass, 2025), AMC23 (AI-MO, 2024b), GSM8K (Cobbe et al., 2021), and OLYMPIADBENCH (He et al., 2024); *scientific reasoning* dataset, GPQA DIAMOND (Rein et al., 2024); *commonsense reasoning* dataset, STRATEGYQA (Geva et al., 2021); and *code reasoning* dataset, LIVECODEBENCH (Jain et al., 2024).

Evaluation metrics. We implement REBALANCE in both Hugging Face Transformers (Wolf et al., 2019) and vLLM (Kwon et al., 2023b). We evaluate using **Pass@1** (\uparrow) and the *average* number of generated tokens **Tok** (\downarrow). Unless otherwise specified, results are reported with the Transformers implementation.

Backbone reasoning models. We conduct experiments on 4 open-source large language models used as backbones: DEEPSPEECH-R1-DISTILL-QWEN (1.5B and 7B) (Guo et al., 2025), QWEN3-14B (Yang et al., 2025a), and QWQ-32B Team (2025). Together, they span 3 model architectures and 4 parameter scales, enabling controlled comparisons across model families and sizes.

Steering extraction and dynamic function fitting. From 500 randomly sampled MATH (Hendrycks et al., 2021) problems, we estimate a *steering vector* and a *control surface* for each backbone once and hold them fixed across all benchmarks.

Baseline methods. We compare REBALANCE against representative training-free methods for efficient inference that do not rely on auxiliary models. (i) **PROMPT-BASED** methods: COD (Xu et al., 2025a) and NOTHINKING (Ma et al., 2025b); (ii) **OUTPUT-BASED** methods: NOWAIT (Wang et al., 2025a); (iii) **DYNAMIC EARLY-EXIT** methods: DYNASOR-COT (Fu et al., 2025), DEER (Yang et al., 2025b), [FLASHTHINK](#) Jiang et al. (2025), and [TRIMR](#) Lin et al. (2025a); (iv) **STEERING** methods: SEAL (Chen et al., 2025) and MANIFOLD STEERING (Huang et al., 2025b). This set covers the major design paradigms (prompting, output-time control, early exiting, and latent steering) under a training-free setting. Further details on additional **PROMPT-BASED** baselines are provided in Appendix G.

Hardware. All experiments were conducted on a single server with **8× NVIDIA RTX PRO 6000 (Blackwell Server Edition)** GPUs.

Decoding settings. Unless otherwise noted, we use nucleus sampling with

`temperature = 0.7, top_p = 0.95, max_generated_tokens = 16000.`

The same decoding configuration is applied across both the Transformers and vLLM backends.

E DETAILS ON BENCHMARKS

MATH-500 (moderate; 500 problems). Comprises 500 problems spanning arithmetic, algebra, geometry, and calculus, with varying difficulty levels. It evaluates a model’s ability in complex mathematical formalism, equation solving, and structured reasoning. (Lightman et al., 2023b)

AIME24 (hard; 30 problems). An Olympiad-style set assessing logical deduction and advanced problem-solving skills; includes official AIME problems from the 2024 cycle. (AI-MO, 2024a)

AIME25 (hard; 30 problems). An updated set from the same AIME competition as AIME24, continuing to target high-level deductive and multi-step mathematical reasoning. (OpenCompass, 2025)

GPQA DIAMOND (hard; 198 problems). A challenging graduate-level subset with multiple-choice questions authored by domain experts in biology, physics, and chemistry. (Rein et al., 2024)

AMC23 (simple; 40 problems). An aggregated 40-problem set based on AMC 12 (2023 A/B). Items are multiple choice and span the standard high-school curriculum (calculus excluded), positioned below AIME-level difficulty. (AI-MO, 2024b)

GSM8K (simple; 1319 problems). Grade-school and middle-school word problems emphasizing short chain-of-thought arithmetic reasoning; commonly used split is $\sim 7.5k$ train / $\sim 1k$ test. (Cobbe et al., 2021)

2160 **OLYMPIADBENCH (hard; 675 problems).** A bilingual math+physics Olympiad-style benchmark
 2161 sourced from international/national olympiads and Gaokao, substantially more challenging than
 2162 standard competition datasets. (He et al., 2024)

2163
 2164 **STRATEGYQA (simple; 2,780 problems).** Yes/No questions requiring *implicit* multi-hop com-
 2165 monsense reasoning; each example provides a decomposition and supporting evidence passages.
 2166 (Geva et al., 2021)

2167
 2168 **LIVECODEBENCH (hard; 400 problems, v1).** A contamination-aware coding benchmark con-
 2169 structed from competitive-programming problems (e.g., LeetCode, AtCoder, Codeforces). Tasks
 2170 require generating runnable programs that are judged by execution-based unit tests, emphasizing
 2171 algorithmic reasoning, data-structure design, and implementation fidelity. We use version v1 with
 2172 400 problems. (Jain et al., 2024)

2173
 2174 F DETAILS ON PROMPTS

2175
 2176 **Math - (MATH-500, AIME 2024, AIME 2025, AMC23, GSM8K, Olympiad).**

2177
 2178 <|System|> Please reason step by step, and place the final answer
 2179 inside \boxed{}.
 2180 <|User|> [question]

2181
 2182 **Science - GPQA.**

2183
 2184 <|System|> Please reason step by step, and place the final answer
 2185 inside \boxed{}.
 2186 <|User|> [question]

2187 Answer with the choice letter only, in \boxed{}. Do not include
 2188 option text.

2189
 2190 **Commonsense - StrategyQA.**

2191
 2192 <|System|> You answer binary commonsense questions. Think step by
 2193 step, then output exactly one final line: \boxed{Yes} or \
 2194 boxed{No}.
 2195 <|User|> [question]

2196 Answer with \boxed{Yes} or \boxed{No} only.

2214

Code - LiveCodeBench.

2215

2216

2217

2218

2219

2220

2221

2222

2223

2224

2225

2226

2227

2228

2229

2230

2231

2232

2233

2234

2235

2236

2237

2238

2239

2240

2241

2242

2243

2244

2245

2246

2247

2248

2249

2250

2251

2252

2253

2254

2255

2256

2257

```

<|User|>
### Instruction: You will be given a question (problem
    specification) and will generate a correct Python program
    that matches the specification and passes all tests. You will
    NOT return anything except for the program.
Question:
[problem]
Ensure that when the python program runs, it reads the inputs,
    runs the algorithm and writes output to STDOUT.
python # YOUR CODE HERE
### Response:<|im_end|><|im_start|>assistant<think>

```

Prompt Used for Steering Vector Extraction in Eq. 3.

2258

2259

2260

2261

2262

2263

2264

2265

2266

2267

2268

2269

2270

2271

2272

2273

2274

2275

2276

2277

2278

2279

2280

2281

2282

2283

2284

2285

2286

2287

2288

2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300

2301

2302

2303

2304

2305

2306

2307

2308

2309

2310

2311

2312

2313

2314

2315

2316

2317

2318

2319

2320

2321

2322

2323

2324

2325

2326

2327

2328

2329

2330

2331

2332

2333

2334

2335

2336

2337

2338

2339

2340

2341

2342

2343

2344

2345

2346

2347

2348

2349

2350

2351

2352

2353

2354

2355

2356

2357

2358

2359

2360

2361

2362

2363

2364

2365

2366

2367

2368

2369

2370

2371

2372

2373

2374

2375

2376

2377

2378

2379

2380

2381

2382

2383

2384

2385

2386

2387

2388

2389

2390

2391

2392

2393

2394

2395

2396

2397

2398

2399

2400

2401

2402

2403

2404

2405

2406

2407

2408

2409

2410

2411

2412

2413

2414

2415

2416

2417

2418

2419

2420

2421

2422

2423

2424

2425

2426

2427

2428

2429

2430

2431

2432

2433

2434

2435

2436

2437

2438

2439

2440

2441

2442

2443

2444

2445

2446

2447

2448

2449

2450

2451

2452

2453

2454

2455

2456

2457

2458

2459

2460

2461

2462

2463

2464

2465

2466

2467

2468

2469

2470

2471

2472

2473

2474

2475

2476

2477

2478

2479

2480

2481

2482

2483

2484

2485

2486

2487

2488

2489

2490

2491

2492

2493

2494

2495

2496

2497

2498

2499

2500

2501

2502

2503

2504

2505

2506

2507

2508

2509

2510

2511

2512

2513

2514

2515

2516

2517

2518

2519

2520

2521

2522

2523

2524

2525

2526

2527

2528

2529

2530

2531

2532

2533

2534

2535

2536

2537

2538

2539

Method	Prompt Modification
CoT (Wei et al., 2022)	<i>Please reason step by step.</i>
CoD (Xu et al., 2025a)	<i>Think step by step, but only keep a minimum draft for each thinking step, with 5 words at most.</i>
CCoT (Renze & Guven, 2024)	<i>Think step by step, Be concise.</i>
CCoT-2-45 (Nayab et al., 2024)	<i>Let's think a bit step by step and limit the answer length to 45 words.</i>
BTC (Ding et al., 2024)	<i>Rapidly evaluate and use the most effective reasoning shortcut to answer the question.</i>
NoThinking (Ma et al., 2025a)	<i><think> Okay, I have finished thinking.</think></i>

Table 23: Prompt modifications used in different reasoning strategies.

has shown that even with the same prompt content, different positioning strategies can significantly affect both the accuracy and efficiency of large language models (Cobbina & Zhou, 2025). However, our main experimental results also indicate that prompt-based methods are not always effective, as the model does not consistently follow instructions. Regarding the broader issue of model controllability, a substantial body of research has already explored the use of reinforcement learning (RL) techniques to achieve more reliable control over large reasoning models (LRMs) (Aggarwal & Welleck, 2025; Yuan et al., 2024). Building upon these insights, exploring how to effectively combine multiple training-free strategies with advanced control methods remains an important direction for future work.

H DETAILED DISCUSSION OF RELATED WORKS

Chain-of-Thought (CoT). CoT prompting elicits intermediate rationales and markedly improves multi-step reasoning (Wei et al., 2022); *self-consistency* further aggregates diverse chains (Wang et al., 2022). Beyond a single chain, search/verification variants under the Tree-of-Thought umbrella include *Tree-of-Thoughts* (Yao et al., 2023), *Stream-of-Search* (Gandhi et al., 2024), *Graph-of-Thoughts* (Besta et al., 2024), *Process Reward Models* (Lightman et al., 2023a), and *RL-based Self-Correction* (Kumar et al., 2024). A converging view is that judiciously increasing *test-time compute*—via multiple paths, search, or verification—can rival or surpass pure parameter scaling for reasoning (Snell et al., 2025).

Latent Reasoning. Latent reasoning shifts the chain-of-thought from discrete tokens to *continuous* hidden representations, reducing tokenized traces and sampling while preserving intermediate signals—thus improving test-time efficiency. Representative approaches include soft thinking with gated hidden-state signals (Zhang et al., 2025d), training models to reason directly in a continuous latent space via internal activations (Hao et al., 2024), compressing long CoT into dense vectors (Cheng & Van Durme, 2024), and assistant-guided soft CoT that maintains a multi-step structure in latent form (Xu et al., 2025b).

Post-Training Methods for Efficient Reasoning. Post-training approaches reduce test-time cost by shaping models’ use of chain-of-thought (CoT) after pretraining. We group them into *SFT-based* and *reinforcement fine-tuning (RFT)-based*

SFT-based. Supervised fine-tuning can induce *conditional brevity*: paired supervision under matched conditions (e.g., long vs. short) teaches when concise reasoning suffices (Kang et al., 2025). A complementary line supervises *compressed* rationales with an explicit compression control at inference; more principled compression schemes further improve faithfulness and controllability (Xia et al., 2025; Yuan et al., 2025).

RFT-based. Reinforcement fine-tuning directly optimizes the accuracy–efficiency trade-off via reward shaping and preference learning. Examples include difficulty-aware ranking to form preference data followed by SimPO optimization (Shen et al., 2025); fixed accuracy–length rewards with PPO (Arora & Zanette, 2025); self-adaptive CoT learning with GRPO (Yang et al., 2025c); and dynami-

2322 cally weighted rewards that balance accuracy and length during PPO training (Su & Cardie, 2025).
 2323 Control-token policies decide *when* to think using GRPO, and explicit reinforcement of *how long* to
 2324 reason enables length control (Aggarwal & Welleck, 2025).

2325 Overall, these methods teach models to reason *when necessary* and remain concise otherwise, im-
 2326 proving the accuracy–latency/token trade-off without increasing parameter count.

2328 **Steering-Based Methods for Efficient Reasoning.** Recent work explores *steering* mechanisms
 2329 that intervene directly in a model’s latent states to improve reasoning efficiency without retraining
 2330 the backbone. Early work has already examined steering in the prompt space to elicit more con-
 2331 trollable model outputs (Liu et al., 2023), and has also leveraged steering-based methods to mitigate
 2332 hallucinations (Liu et al., 2024). Recent work has extensively explored using steering-based methods
 2333 to enable more efficient reasoning. SEAL (Chen et al., 2025) partitions model thoughts into *exe-*
 2334 *cution*, *reflection*, and *transition* phases, and uses a small amount of training data to bias the model
 2335 toward the *execution* mode. MANIFOLD STEERING (Huang et al., 2025b) constructs *redundant* ver-
 2336 *versus* *concise* datasets based on response length and keyword density, derives a steering vector from
 2337 them, and applies it to *all tokens across all layers*. REASONING STRENGTH PLANNING (Sheng
 2338 et al., 2025) further introduces a *pre-allocated direction vector* injected into the activation corre-
 2339 *spending* to the *<think>* token, whose magnitude encodes the desired reasoning strength in terms
 2340 of the target number of reasoning tokens, with steering consistently applied at each layer for every
 2341 generated token. Unlike the static steering methods above, CONTROLLING THINKING SPEED(Lin
 2342 et al., 2025b) introduces a dynamic variant. They construct steering vectors by pairing long and short
 2343 correct responses and extracting the hidden states of the last token in the first two reasoning steps
 2344 at a chosen layer. A sliding-window controller then adjusts the steering strength based on token-
 2345 level difficulty, measured via the Jensen–Shannon divergence between shallow- and deep-layer logit
 2346 distributions, decreasing or increasing the magnitude according to a window-specific threshold.

2347 However, the adjustment mechanism remains fundamentally one-directional in how steering mag-
 2348 nitude is updated. Moreover, the steering-based methods above focus primarily on mitigating over-
 2349 thinking, yet overlook a complementary issue: alleviating overthinking often introduces or amplifies
 2350 underthinking (Wang et al., 2025c). REBALANCE addresses this gap through a bidirectional dynamic
 2351 steering mechanism that uses real-time confidence during reasoning as an indicator to assess the ten-
 2352 dency toward overthinking or underthinking, dynamically adjusting both the direction and intensity
 2353 of steering to achieve efficient reasoning with balanced thinking.

2354 **Early-Exit Methods for Efficient Reasoning.** A prominent line of work toward efficient reasoning
 2355 involves enabling models to exit early from their reasoning process once sufficient evidence for
 2356 an answer has been gathered. Representative approaches such as TRIMR (Lin et al., 2025a) and
 2357 FLASHTHINK (Jiang et al., 2025) employ an external instruction-following model to monitor the
 2358 target model’s chain-of-thought: when the monitor deems further reasoning unnecessary, it halts the
 2359 process and triggers answer generation. Other methods, including DEER (Yang et al., 2025b) and
 2360 DYNASOR-COT (Fu et al., 2025), instead rely on internal signals, such as the model’s confidence
 2361 or entropy over candidate answers, to determine an appropriate exit point. While these strategies
 2362 effectively reduce token consumption by forcibly terminating redundant reasoning, they all share a
 2363 fundamental limitation: the decision to terminate is binary and coarse-grained, typically applied at
 2364 the level of the entire reasoning path (e.g., sub-thoughts in TRIMR or reasoning chunks in FLASH-
 2365 THINK). This rigid binary selection risks discarding potentially valuable reasoning steps, thereby
 2366 inducing additional underthinking (Wang et al., 2025c). Although TRIMR includes certain mech-
 2367 anisms addressing underthinking, it primarily opts for abandoning reasoning upon detecting under-
 2368 thinking, which leans towards engineering-driven token length optimization rather than genuinely
 2369 addressing underthinking itself. Moreover, both TRIMR and FLASHTHINK depend on handcrafted
 2370 keyword triggers to identify termination points, limiting their adaptability across models and tasks.

2371 In contrast, REBALANCE departs fundamentally from this paradigm. Rather than merely mitigating
 2372 overthinking through early exit, ReBalance is explicitly designed to simultaneously mitigate over-
 2373 thinking and prevent underthinking. It achieves this not by discarding reasoning paths, but by lever-
 2374 aging confidence as an indicator to dynamically detect the model’s tendency toward overthinking
 2375 or underthinking in real-time. Based on this detection, it adaptively adjusts the strength and direc-
 2376 tion of steering, thereby dynamically controlling the model’s behavior to keep its reasoning state
 2377 consistently within the reasoning boundary. This enables balanced thinking without requiring any

2376 additional verifiers or inference stages, making ReBalance an efficient, dynamic, and fine-grained
 2377 reasoning acceleration approach.
 2378

2379 **Acceleration for Efficient Reasoning.** Beyond training-time efficiency, a complementary line of
 2380 work targets *inference-time* acceleration—reducing latency and improving throughput under fixed
 2381 hardware budgets. One class of methods exploits *speculative decoding*, drafting tokens with a
 2382 lightweight proposer and verifying them with the target model to amortize compute (Leviathan et al.,
 2383 2023). A second class reduces memory and scheduling overhead via *paged KV-cache* management
 2384 and continuous batching, enabling high-utilization serving at scale (Kwon et al., 2023a). A third line
 2385 optimizes the *sampling process* itself: Monte Carlo tree or search-style strategies for data synthesis
 2386 (Li et al., 2025), best-of- n reasoning accelerated by speculative rejection (Sun et al., 2024), and
 2387 early-decoding schemes that *self-estimate* the necessary n to balance quality and cost (Wang et al.,
 2388 2025b).
 2389

2390 **Small Language Models (SLMs) for Efficient Reasoning.** A complementary line of work pur-
 2391 sues efficient inference by *compressing* or *transferring* reasoning ability into smaller backbones.
 2392 First, *quantization* can degrade multi-step reasoning if applied naively, calling for calibration-aware
 2393 schemes and mixed-precision designs (Liu et al., 2025b). Second, *distillation* transfers long-horizon
 2394 reasoning into compact policies by (i) shortening and regularizing chain-of-thought traces and (ii)
 2395 internalizing deliberate reasoning into fast feedforward behavior (Luo et al., 2025; Yu et al., 2024).
 2396 Third, *pruning/compression* benchmarks sparsity and related compression knobs on complex rea-
 2397 soning tasks, revealing sensitivity to where and how compression is applied (Zhang et al., 2025b).
 2398 Finally, a recent assessment reports the combined effects of distillation and pruning on SLMs, high-
 2399 lighting regimes where small models recover strong reasoning at a fraction of the compute (Srivastava
 2400 et al., 2025).
 2401

I EFFICIENCY ANALYSIS

2404 We conduct a comprehensive efficiency evaluation of ReBalance by comparing its inference over-
 2405 head against both the baseline and other efficient reasoning methods. Specifically, we report four
 2406 key metrics: tokens per second (TPS), time per request (TPR), and additional GPU memory con-
 2407 sumption relative to the baseline.
 2408

2409 Specifically, to ensure a thorough comparison, we include prompt-based methods (NoThinking,
 2410 CoD), early-exit methods (FlashThink, TrimR), and latent steering methods (SEAL). Different from
 2411 existing methods, ReBalance preserves an effective thinking process by promoting exploration when
 2412 necessary to avoid underthinking, particularly when the model faces difficult reasoning problems.
 2413 As a result, it is more likely to incur efficiency overheads on challenging datasets. In the following,
 2414 we adopt AIME24 for the efficiency evaluation, as shown in Tab. 24.

- 2415 • **Token generation efficiency:** By observing TPS, it can be seen that since the confidence
 2416 utilized by ReBalance can be directly obtained from the log probability of each token’s
 2417 decoded output, and the dynamic function introduced in this method is very lightweight, the
 2418 proposed dynamic adjustment logic has a negligible impact on the single-token generation
 2419 time compared to the baseline.
- 2420 • **Reasoning time acceleration:** As shown by TPS and #Tokens, ReBalance significantly
 2421 shortens reasoning length without compromising token generation efficiency, yielding **1.5**×
 2422 and **1.4**× TPR speedups over DeepSeek-R1-Distill-Qwen-7B and QwQ-32B, respectively.
 2423
- 2424 • **Additional GPU memory usage:** Although ReBalance requires the use of extracted steer-
 2425 ing vectors during reasoning, their GPU memory footprint is minimal (e.g., the vector
 2426 size for QwQ-32B is only 22 KB). In contrast, early-exit methods, such as FlashThink and
 2427 TrimR, require additional verifiers (usually at least 7B), introducing extra GPU memory
 2428 usage and communication load.

2429 In conclusion, ReBalance achieves outstanding performance in tokens per second, time per request,
 2430 additional GPU memory usage, and generated sequence length.

Method	TPS	TPR (s)	Additional GPU Memory (GB)	#Tokens
DeepSeek-R1-Distill-Qwen-7B				
Baseline	80.2	174.6	—	13994
NoThinking	80.1	55.2	0.0	4427
CoD	80.2	145.5	0.0	11663
FlashThink	73.8	135.9	18.3	10034
TrimR	48.8	147.7	15.4	7213
SEAL	78.8	128.3	0.0	10112
ReBalance (Ours)	78.5	114.8	0.0	9012
QwQ-32B				
Baseline	20.5	698.6	—	14342
NoThinking	20.6	509.6	0.0	10507
CoD	20.7	552.0	0.0	11438
FlashThink	14.9	673.4	18.3	10034
TrimR	6.5	1289.8	15.7	8345
SEAL	20.3	508.8	0.0	10344
ReBalance (Ours)	20.5	504.4	0.0	10350

Table 24: Efficiency analysis on DeepSeek-R1-Distill-Qwen-7B and QwQ-32B.

J THE USE OF LARGE LANGUAGE MODELS

In this work, large language models are used exclusively for polishing the writing and checking grammar. They are not involved in research ideation, experimental design, data analysis, or the formulation of conclusions. All substantive intellectual contributions are made by the authors.

K ETHICS STATEMENT

We confirm that this research adheres to the ICLR Code of Ethics. Our method is purely algorithmic and theoretical in nature, involving no human subjects, personal data, or real-world deployment that could lead to harmful, biased, or discriminatory outcomes. We have carefully considered potential ethical implications—including fairness, transparency, and societal impact—and conclude that the proposed approach poses no significant ethical concerns. All experiments were conducted on publicly available datasets.

L REPRODUCIBILITY STATEMENT

To facilitate replication, we fix the random seed to **42** for all sampling and evaluation. All experiments are run on the same hardware ($8 \times$ NVIDIA RTX PRO 6000 GPUs) with unified decoding configurations (`temperature=0.7`, `top_p=0.95`, `max_generated_tokens=16000`). Steering vectors and control surfaces are extracted once from a random sample of 500 MATH problems and then reused across all benchmarks. Together with standardized evaluation metrics (Pass@1 and average token length), these settings ensure that our results can be reproduced reliably.

2484 M CASE STUDY
2485
2486
2487
2488
2489
2490
2491

2492 Question:

2493 For what real values of x is $-4 < x^4 + 4x^2 < 21$ satisfied?
2494 Express your answer in interval notation.2495 Answer: $(-\sqrt{3}, \sqrt{3})$

Generated Responses:<think>\nAlright, so I have this inequality to solve: $-4 < x^4 + 4x^2 < 21$. Hmm, okay. It looks like a quartic inequality because of the x^4 term. I remember that to solve such inequalities, I can break it down into two separate inequalities and solve each part separately. That is, first solve $-4 < x^4 + 4x^2$ and then solve $x^4 + 4x^2 < 21$. After that, I can find the intersection of the solutions from both parts to get the final answer.\n\nLet me start by rewriting the inequality:
 $-4 < x^4 + 4x^2 < 21$. So, I'll split this into two parts:
 $x^4 + 4x^2 > -4$ and $x^4 + 4x^2 < 21$. Starting with the first inequality: $x^4 + 4x^2 > -4$. Hmm, x^4 is always non-negative since any real number raised to an even power is positive or zero. Similarly, $4x^2$ is also non-negative for the same reason. So, adding them together, $x^4 + 4x^2$ is always greater than or equal to zero. But here, the right side is -4 , which is negative. So, is it possible for $x^4 + 4x^2$ to be greater than -4 ? Well, since $x^4 + 4x^2$ is always at least zero, and zero is greater than -4 , that inequality $x^4 + 4x^2 > -4$ is always true for all real numbers x . So, the first part doesn't impose any restrictions on x . That's a relief; I don't have to worry about excluding any x from the solution set because the left inequality is always satisfied.\n\nNow, moving on to the second inequality: $x^4 + 4x^2 < 21$. This one is more challenging. Let's try to simplify or rearrange it. Maybe I can set $y = x^2$ to make it easier to handle. Since x^2 is always non-negative, $y \geq 0$. Substituting, the inequality becomes $y^2 + 4y < 21$. So, let me write that out:
 $y^2 + 4y - 21 < 0$. Now, I can solve this quadratic inequality in terms of y . First, let's find the roots of the equation $y^2 + 4y - 21 = 0$. Using the quadratic formula: $y = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$. Here, $a = 1$, $b = 4$, $c = -21$. Plugging in, we get:
 $y = \frac{-4 \pm \sqrt{(16 - 4 \cdot 1 \cdot -21)}}{2} = \frac{-4 \pm \sqrt{16 + 84}}{2} = \frac{-4 \pm \sqrt{100}}{2} = \frac{-4 \pm 10}{2} = 3, -7$. So, the roots are: $y = 3$ and $y = -7$. But since $y = x^2$ and x^2 can't be negative, y can't be -7 . So, the only relevant root is $y = 3$. Now, to solve the inequality $y^2 + 4y - 21 < 0$, we can analyze the intervals determined by the roots. Since the quadratic opens upwards (because the coefficient of y^2 is positive), the quadratic will be below zero between its two roots. However, one of the roots is negative ($y = -7$), which we can ignore because y is non-negative. So, the relevant interval is from $y = -7$ to $y = 3$. But since y can't be negative, we only consider y from 0 to 3. Thus, the inequality $y^2 + 4y - 21 < 0$ holds when y is between 0 and 3, i.e., $0 < y < 3$. But remember, $y = x^2$, so substituting back, we get:
 $0 < x^2 < 3$. This inequality implies that x^2 is between 0 and 3. To find x , we can take square roots. First, since $x^2 > 0$, that means $x \neq 0$. But we have to consider $x^2 < 3$ as well. So, $x^2 < 3$ implies that $|x| < \sqrt{3}$, which means x is in the interval $(-\sqrt{3}, \sqrt{3})$. Putting it all together, the second inequality $x^4 + 4x^2 < 21$ is satisfied when x is between $-\sqrt{3}$ and $\sqrt{3}$, excluding zero? Wait, hold on. No, actually, since $x^2 > 0$, the inequality $x^2 < 3$ holds for all x in $(-\sqrt{3}, \sqrt{3})$ except $x = 0$. But actually, when $x = 0$, $x^4 + 4x^2 = 0$, which is less than 21. So, actually, $x = 0$ is included. Wait, let's think again! When we transformed $x^4 + 4x^2 < 21$ into $y^2 + 4y - 21 < 0$, we found that y must be between 0 and 3. Since $y = x^2$, that translates to x^2 between 0 and 3. So, $x^2 < 3$, which is equivalent to $|x| < \sqrt{3}$, meaning x is in $(-\sqrt{3}, \sqrt{3})$. So, actually, x can be zero here because when x is zero, $x^4 + 4x^2 = 0$, which is indeed less than 21. So, the interval is from $-\sqrt{3}$ to $\sqrt{3}$, including zero. Wait, but in the first inequality, $x^4 + 4x^2 > -4$ is always true. So, combining both inequalities, the solution is all x such that x is in $(-\sqrt{3}, \sqrt{3})$. Is that right? Wait, but when x is zero, $x^4 + 4x^2$ is zero, which is less than 21, so it's included. So, the solution set is all real numbers x such that x is between $-\sqrt{3}$ and $\sqrt{3}$. But let me double-check this. Let's test a value inside the interval, say $x = 1$. Then, $x^4 + 4x^2 = 1 + 4 = 5$, which is between -4 and 21. So, it's valid. What about $x = \sqrt{3}$? Then, $x^4 + 4x^2 = (\sqrt{3})^4 + 4(\sqrt{3})^2 = 9 + 12 = 21$, which is equal to the upper limit. So, $x = \sqrt{3}$ is not included because the inequality is strict ($<$). Similarly, $x = -\sqrt{3}$ is also not included. What about $x = 2$? Then, $x^4 + 4x^2 = 16 + 16 = 32$, which is greater than 21, so it's excluded. So, indeed, the solution is x between $-\sqrt{3}$ and $\sqrt{3}$, not including the endpoints. Wait, but let me check $x = 0.5$, which is inside the interval. $x^4 + 4x^2 = (0.5)^4 + 4(0.5)^2 = 0.0625 + 4 \cdot 0.25 = 0.0625 + 1 = 1.0625$, which is between -4 and 21, so it's valid. What about $x = 1.5$, which is still less than $\sqrt{3}$ (since $\sqrt{3} \approx 1.732$). $x^4 + 4x^2 = (1.5)^4 + 4(1.5)^2 = 5.0625 + 4 \cdot 2.25 = 5.0625 + 9 = 14.0625$, which is still less than 21, so it's valid. How about $x = \sqrt{3} \approx 1.732$? As we saw earlier, $x^4 + 4x^2 = 21$, so it's exactly 21, which isn't included. Similarly, $x = -\sqrt{3}$ gives 21 as well. So, putting it all together, the solution is all real numbers x such that $-\sqrt{3} < x < \sqrt{3}$. Wait, let me think again about the first inequality. Since $x^4 + 4x^2$ is always greater than or equal to zero, which is greater than -4, so the first inequality doesn't restrict x at all. So, the only restriction comes from the second inequality. Therefore, combining both inequalities, the solution is $x \in (-\sqrt{3}, \sqrt{3})$. But just to make sure, let me consider if there are any points outside of this interval that satisfy both inequalities. For example, $x = 2$, which is outside, as we saw, gives $x^4 + 4x^2 = 32$, which is greater than 21, so it's excluded. Similarly, $x = -2$ gives the same result. What about $x = \sqrt{3}$, as we saw, gives 21, which is excluded. $x = 0$, gives 0, which is included. So, yeah, it seems that the only x that satisfy both inequalities are those in the interval $(-\sqrt{3}, \sqrt{3})$. Therefore, in interval notation, that would be $(-\sqrt{3}, \sqrt{3})$. Wait, just to double-check, is there any way for $x^4 + 4x^2$ to be less than or equal to -4? Since $x^4 + 4x^2$ is always non-negative, it can't be less than -4. So, the only time when it's less than 21 is when it's between 0 and 21. But $x^4 + 4x^2$ is 0 at $x = 0$, and as x increases or decreases, it grows. So, when does it reach 21? At $x = \pm\sqrt{3}$, as we saw. Therefore, the solution is indeed $x \in (-\sqrt{3}, \sqrt{3})$. So, in interval notation, that's written as $(-\sqrt{3}, \sqrt{3})$. But wait, just to make sure, let me consider the quadratic in y . So, we had $y^2 + 4y - 21 < 0$, with $y = x^2 \geq 0$. We found that y is between 0 and 3, which corresponds to x between $-\sqrt{3}$ and $\sqrt{3}$. Yes, that seems consistent. So, I think I can confidently say that the solution is all real numbers x such that x is between $-\sqrt{3}$ and $\sqrt{3}$. Final Answer: The solution is $(-\sqrt{3}, \sqrt{3})$.</think>.....Therefore, the solution to the inequality $(-4 < x^4 + 4x^2 < 21)$ is: $(-\sqrt{3}, \sqrt{3})$.

2530 Correct Answer, But answer very long. (1659 tokens)
2531
2532

2533 **Figure 14:** A DeepSeek-R1-Distill-Qwen-1.5B inference case shows that, although the model arrives at the
2534 correct solution, its reasoning is rife with redundancy and repetition. Beyond the necessary boundary checks,
2535 it inspects numerous inconsequential points, incurring substantial token overhead.

2538

2539

2540

2541

2542

2543

2544

2545

2546

Question:

For what real values of x is $-4 < x^4 + 4x^2 < 21$ satisfied?
Express your answer in interval notation.Answer: $(-\sqrt{3}, \sqrt{3})$

2547

2548

2549

2550

2551

2552

2553

2554

2555

Generated Responses: "Okay, so I have this inequality to solve: $-4 < x^4 + 4x^2 < 21$. Hmm, I need to find all real numbers x that satisfy this. Let me think about how to approach this.\n\nFirst, I notice that the expression is a quartic polynomial, but it's only in terms of x^2 . Maybe I can make a substitution to simplify it. Let me set $y = x^2$. Then, the inequality becomes $-4 < y^2 + 4y < 21$. That seems easier to handle because it's a quadratic in terms of y .\n\nSo, now I have two inequalities to solve:\n\n1. $y^2 + 4y > -4$ \n2. $y^2 + 4y < 21$ \n\nLet me tackle the first one: $y^2 + 4y > -4$. I can rewrite this as $y^2 + 4y + 4 > 0$. Hmm, that looks like a perfect square. Let me check: $(y + 2)^2 = y^2 + 4y + 4$. Yes, that's correct. So, the inequality becomes $(y + 2)^2 > 0$. When is a square greater than zero? Well, squares are always non-negative, so $(y + 2)^2$ is greater than zero when $y + 2$ is not equal to zero. That means $y \neq -2$. But wait, y was defined as x^2 , and x^2 is always non-negative. So, $y = -2$ is not possible because x^2 can't be negative. So, the inequality $(y + 2)^2 > 0$ is always true except when $y = -2$, which isn't in our domain. So, the first inequality is always satisfied for all real x except when $y = -2$, but since y can't be negative, this inequality doesn't restrict x at all. So, the first part doesn't impose any restrictions.\n\nNow, moving on to the second inequality: $y^2 + 4y < 21$. Let me rewrite this as $y^2 + 4y - 21 < 0$. To solve this quadratic inequality, I need to find the roots of the equation $y^2 + 4y - 21 = 0$. I can use the quadratic formula: $y = [-b \pm \sqrt{b^2 - 4ac}] / (2a)$. Here, $a = 1$, $b = 4$, $c = -21$. Calculating the discriminant: $b^2 - 4ac = 16 - 4(1)(-21) = 16 + 84 = 100$. That's a perfect square, which is nice. So, the roots are $y = [-4 \pm \sqrt{100}] / 2 = [-4 \pm 10] / 2$. Calculating the two roots:\n\n1. $y = (-4 + 10)/2 = 6/2 = 3$ \n2. $y = (-4 - 10)/2 = -14/2 = -7$ \n\nSo, the roots are $y = 3$ and $y = -7$. Since $y = x^2$, and x^2 can't be negative, y can't be -7 . Therefore, the relevant root is $y = 3$. Now, the quadratic $y^2 + 4y - 21$ factors as $(y - 3)(y + 7)$. But since y can't be negative, we only consider $y \geq 0$. So, the quadratic is negative between its roots. But since one root is negative and the other is positive, the quadratic is negative between $y = -7$ and $y = 3$. However, since y can't be less than 0, the relevant interval is y between 0 and 3.\n\nTherefore, the inequality $y^2 + 4y - 21 < 0$ holds when $0 < y < 3$. Since $y = x^2$, this translates to $0 < x^2 < 3$. Now, let's solve $0 < x^2 < 3$. The inequality $0 < x^2$ means that x is not zero, so $x \neq 0$. The inequality $x^2 < 3$ implies that $-\sqrt{3} < x < \sqrt{3}$. Combining these, we get $-\sqrt{3} < x < 0$ and $0 < x < \sqrt{3}$. So, in interval notation, that's $(-\sqrt{3}, 0) \cup (0, \sqrt{3})$. But wait, let me double-check. The original inequality was $-4 < x^4 + 4x^2 < 21$. We transformed it to $y^2 + 4y < 21$, which gave us $0 < y < 3$, so $0 < x^2 < 3$, which is indeed $-\sqrt{3} < x < \sqrt{3}$ excluding $x=0$. So, that seems correct.\n\nBut just to be thorough, let me test some values to make sure.\n\nFirst, let's test $x = 1$. Then, $x^4 + 4x^2 = 1 + 4 = 5$. Is 5 between -4 and 21? Yes, it is.\n\nWhat about $x = \sqrt{3}$? Then, $x^4 + 4x^2 = (\sqrt{3})^4 + 4*(\sqrt{3})^2 = (9) + 4*(3) = 9 + 12 = 21$. So, 21 is the upper bound, but the inequality is strict, so $x = \sqrt{3}$ is not included. Similarly, $x = -\sqrt{3}$ is not included.\n\nWhat about $x = 0$? Then, $x^4 + 4x^2 = 0 + 0 = 0$, which is between -4 and 21. So, $x = 0$ is included?\n\nThe solution involves substituting $y = x^2$, transforming the inequality into a quadratic, and solving for y . The resulting interval for y is $0 < y < 3$, which translates to x being in $(-\sqrt{3}, 0) \cup (0, \sqrt{3})$.** Answer:** The solution in interval notation is $(-\sqrt{3}, 0) \cup (0, \sqrt{3})$.

2576

2577

2578

2579

2580

2581

2582

2583

2584

2585

2586

Figure 15: In a DeepSeek-R1-Distill-Qwen-1.5B inference case, applying existing overthinking-mitigation techniques reduces token usage relative to the baseline; however, the absence of verification steps results in an incorrect answer.

2587

2588

2589

2590

2591

UnderThinking, Should Double-Check

Answer short, But incorrect Answer (834 tokens)

2592

2593

2594

2595

2596

2597

2598

2599

2600

2601

Question:

For what real values of x is $-4 < x^4 + 4x^2 < 21$ satisfied?

Express your answer in interval notation.

Answer: $(-\sqrt{3}, \sqrt{3})$

2603

2604

2605

2606

2607

2608

Generated Responses: "Okay, so I have this inequality to solve: $-4 < x^4 + 4x^2 < 21$. Hmm, it's a quartic inequality because of the x^4 term. I remember that for inequalities with polynomials, especially even degree ones, the graphs tend to have a U-shape on both ends. So, the function $f(x) = x^4 + 4x^2$ will go to infinity as x approaches both positive and negative infinity. That means the function will have a minimum somewhere in the middle. First, maybe I should try to solve the inequality step by step. It's like two separate inequalities: $x^4 + 4x^2 > -4$ and $x^4 + 4x^2 < 21$. But since x^4 and x^2 are always non-negative (because any real number squared is non-negative, and raising it to the fourth power is also non-negative), their sum will also be non-negative. So, $x^4 + 4x^2$ is always greater than or equal to zero. That means the first inequality, $x^4 + 4x^2 > -4$, is always true because a non-negative number is always greater than -4. So, I don't have to worry about that part; it doesn't restrict the solution set. So, the real work is solving the second inequality: $x^4 + 4x^2 < 21$. Let me rewrite that as $x^4 + 4x^2 - 21 < 0$. Maybe I can factor this expression or find its roots to figure out where it's negative. Let me try to factor $x^4 + 4x^2 - 21$. Hmm, noticing that it's a quadratic in terms of x^2 , I can let $y = x^2$, so the equation becomes $y^2 + 4y - 21 = 0$. Let me solve for y . Using the quadratic formula: $y = [-4 \pm \sqrt{16 + 84}] / 2 = [-4 \pm \sqrt{100}] / 2 = [-4 \pm 10] / 2$. So, the solutions are $y = (-4 + 10) / 2 = 6 / 2 = 3$, and $y = (-4 - 10) / 2 = -14 / 2 = -7$. So, $y = 3$ or $y = -7$. But $y = x^2$, which is always non-negative, so $y = -7$ doesn't make sense in this context. So, the only relevant solution is $y = 3$. Therefore, $x^2 = 3$, so $x = \sqrt{3}$ or $x = -\sqrt{3}$. So, that means the quartic function $x^4 + 4x^2 - 21$ is zero at $x = \sqrt{3}$ and $x = -\sqrt{3}$. Since it's a quartic function opening upwards (because the coefficient of x^4 is positive), the graph will touch the x-axis at $x = \sqrt{3}$ and $x = -\sqrt{3}$, and it will be below the x-axis between these two roots. So, the inequality $x^4 + 4x^2 - 21 < 0$ holds when x is between $-\sqrt{3}$ and $\sqrt{3}$. Therefore, the solution to the original inequality $-4 < x^4 + 4x^2 < 21$ is the same as $x^4 + 4x^2 < 21$, which is when x is between $-\sqrt{3}$ and $\sqrt{3}$. But wait, I should check if there are any other intervals or points where the inequality might hold. Since the quartic is always greater than or equal to zero, and we're looking for it to be less than 21, which is a positive number, so it's only between the two roots where it dips below 21. So, that means x must satisfy $-\sqrt{3} < x < \sqrt{3}$. Let me verify that with some test points. For $x = 0: 0^4 + 4*(0)^2 = 0$, which is between -4 and 21, so that's good. For $x = 1: 1 + 4 = 5$, which is between -4 and 21, so that's fine. For $x = 2: 16 + 16 = 32$, which is greater than 21, so that's outside the solution. For $x = \sqrt{3}$: approx 1.732, the function equals 21, so that's the boundary. Similarly, for $x = -\sqrt{3}$: same result. So, the interval from $-\sqrt{3}$ to $\sqrt{3}$ is where the function is less than 21, and since the other inequality is always true, the solution is just that interval. Therefore, in interval notation, the solution is $(-\sqrt{3}, \sqrt{3})$. Final Answer** The solution is $(-\sqrt{3}, \sqrt{3})$. To solve the inequality $(-4 < x^4 + 4x^2 < 21)$, we start by noting that $(x^4 + 4x^2)$ is always non-negative because it is a sum of squares. Therefore, the inequality $(x^4 + 4x^2 > -4)$ is always true. Next, we focus on the inequality $(x^4 + 4x^2 < 21)$. We rewrite this as $(x^4 + 4x^2 - 21 < 0)$. Letting $(y = x^2)$, we transform the equation into a quadratic: $(y^2 + 4y - 21 = 0)$. Solving this quadratic equation using the quadratic formula, we find the roots: $y = \frac{-4 \pm \sqrt{16 + 84}}{2} = \frac{-4 \pm \sqrt{100}}{2} = \frac{-4 \pm 10}{2}$. This gives us $y = 3$ and $y = -7$. Since $y = x^2$ must be non-negative, we discard $y = -7$. Thus, $y = 3$ gives $x = \sqrt{3}$ or $x = -\sqrt{3}$. The function $(x^4 + 4x^2 - 21)$ is a quartic that opens upwards and is zero at $x = \sqrt{3}$ and $x = -\sqrt{3}$. It is negative between these roots. Therefore, the inequality $(x^4 + 4x^2 < 21)$ holds for x in the interval $(-\sqrt{3}, \sqrt{3})$. Thus, the solution to the original inequality $(-4 < x^4 + 4x^2 < 21)$ is:

Efficiently Check the Boundary.

Answer short and Correct (875 tokens)

2635

2636

2637

2638

2639

2640

2641

Figure 16: In a DeepSeek-R1-Distill-Qwen-1.5B inference example, the incorporation of judicious verification steps yields a correct and succinct response.

2643

2644

2645