

# Think How to Think: Mitigating Overthinking with Autonomous Difficulty Cognition in Large Reasoning Models

Yongjiang Liu<sup>1</sup>, Haoxi Li<sup>1</sup>, Xiaosong Ma<sup>2</sup>, Jie Zhang<sup>1\*</sup>, Song Guo<sup>1</sup>,

<sup>1</sup>The Hong Kong University of Science and Technology

<sup>2</sup>The Hong Kong Polytechnic University

{yliunr, hligb}@connect.ust.hk

{csejzhang, songguo}@cse.ust.hk

maxiaosong16@gmail.com

## Abstract

Recent Large Reasoning Models (LRMs) excel at complex reasoning tasks but often suffer from overthinking, generating overly long and redundant reasoning trajectories. To explore its essence, our empirical analysis reveals that LRMs are primarily limited to recognizing task properties (i.e., difficulty levels) like humans before solving the problem, leading to a one-size-fits-all reasoning strategy. This observation motivates a fundamental question: *Can we explicitly bootstrap such ability to alleviate overthinking in LRMs?* To this end, we propose **Think-How-to-Think (TH2T)**, a novel two-stage fine-tuning strategy that progressively inspires LRMs’ difficulty cognition and redundancy cognition of LRMs. Specifically, we first inject **Difficulty Dypnosis** into output prefixes as cues for global, prospective reasoning strategy selection, stimulating the model’s sharper sensitivity to task complexity and adaptive control of reasoning depth. Then, we incorporate **Redundancy Hypnosis** into in-progress reasoning steps, which serve as local, retrospective signals for behavior correction by identifying and eliminating superfluous reasoning detours. Experiments across 7B/14B/32B models demonstrate that **TH2T** significantly reduces inference costs by over 70% on easy tasks and 40% on complex ones without compromising performance. The resultant models exhibit a nascent ability for difficulty-aware reasoning, effectively mitigating behaviors like excessive reflection and looping, thereby paving the way for more cognitively efficient LRMs.

## 1 Introduction

The emergence of Large Reasoning Models (LRMs) has led to substantial improvements in problem-solving capabilities (Sun et al., 2024b; Yu et al., 2024; Team et al., 2025), particularly through the use of elaborate and sequential reasoning enabled by test-time scaling (Snell et al., 2025). Many

\*Corresponding author.

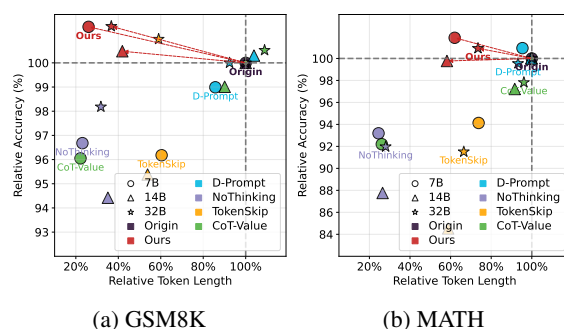


Figure 1: Overall comparison with baselines on R1-Distill-Qwen series models. Relative accuracy and token length of original LRM are set as 100%.

models, such as OpenAI-O1 (OpenAI et al., 2024a), Deepseek-R1 (Guo et al., 2025), QwQ (Team, 2025) leverage long Chain-of-Thought (Wei et al., 2022; Wang et al., 2023) to approach complex tasks step-by-step, enabling a human-like thinking pattern with reflection, backtracking, and self-validation (Xie et al., 2025; Liu et al., 2025). This reasoning paradigm demonstrably enhances the models’ problem-solving capabilities and yields promising results. However, lengthy responses introduce the problem of “overthinking” (Chen et al., 2025b), wherein model performance initially improves with extended reasoning chains but deteriorates beyond a certain length. This inefficiency is starkly evident when trivial queries, such as “*What is 1 plus 1?*”, trigger disproportionately complex reasoning paths. Such redundant deliberation not only inflates computational cost and latency but also degrades the user experience. These lead us to a central inquiry: *Do LRMs possess the cognitive judgment to dynamically regulate their reasoning depth in response to varying problem difficulty?*

Inspiration for a solution lies in the human cognitive science: the dual-process theory of human reasoning. This theory posits a cognitive architecture composed of two distinct but complementary systems: a fast, concise, and intuitive *System 1*, and

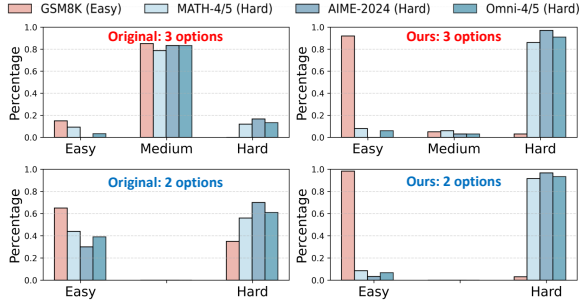


Figure 2: LRM’s difficulty cognition. LRM is prompted to assess the difficulty level of questions from the relatively straightforward GSM8K (Cobbe et al., 2021) and highly complex tasks (MATH (Hendrycks et al., 2021b), AIME2024 and OmniMath (Gao et al., 2024)). Our method substantially reduces difficulty cognition conflation, consistently across both 3-option (e.g., Easy, Medium, Hard) and 2-option (e.g., Easy, Hard) setups.

a slow, deliberate, and analytical *System 2* (Kahneman, 2011; Evans and Stanovich, 2013). Before engaging in problem-solving, humans often perform an initial assessment of a task’s difficulty, allowing them to allocate cognitive resources appropriately between these two systems (Efklides, 2011; Flavell, 1979). For straightforward tasks, reliance on *System 1* suffices, whereas complex problems necessitate the engagement of *System 2*. Our empirical analysis (Fig. 2) reveals that they approach disparate tasks—from the straightforward (e.g., GSM8K) to the highly complex (e.g., AIME2024)—with a monolithic and inefficient reasoning strategy. These limitations highlight the need for LRMs to develop autonomous difficulty cognition—to “think how to think”.

In this paper, we propose **Think-How-to-Think (TH2T)**, a novel two-stage fine-tuning strategy that progressively inspires LRMs difficulty cognition and redundancy cognition. In contrast to prompt-based input guidance for instruction adherence, our approach emphasizes output-level interventions, directly intervening inside the reasoning trajectory to establish a self-hypnosis mechanism. In the first stage, we design **Difficulty Hypnosis (DH)** to stimulate the model’s awareness of task difficulty. DH introduces difficulty cues into output prefixes, which act as guiding signals for global, prospective strategy selection. With DH supervision from a hybrid dataset featuring both concise and elaborate reasoning trajectories, the model calibrates its sensitivity to task difficulty and adaptive control of reasoning depth, ultimately curtailing the impulse to overthink on straightforward prob-

lems through structured cognitive self-guidance. In the second stage, we design **Redundancy Hypnosis (RH)** to cultivate the model’s awareness of redundant reasoning. RH specifically targets fine-grained inference patterns such as redundant yet logically complete structures (e.g., reflections or double-checks) and terminal reasoning loops frequently observed in complex task generation. By embedding redundancy cues into these in-progress structures, RH provides local, retrospective signals for behavior correction, enabling the model to skip or truncate superfluous steps. In this way, the model strengthens its redundancy cognition and reduces superfluous reasoning detours. Overall, our contributions can be outlined as follows:

- We reveal that current LRMs fall short of perceiving task complexity. To mitigate overthinking, we propose enhancing LRMs with the ability to cognize and adapt to varying levels of problem difficulty autonomously.
- To align with human *System 1* and *System 2* thinking patterns, we propose a novel two-stage fine-tuning strategy with a self-hypnosis mechanism for prospective and retrospective intervention, thus progressively inspiring LRMs’ difficulty and redundancy cognition.
- Experiments demonstrate over 70% length reduction on easy tasks and 40% on hard tasks with maintained robust performance, as shown in Fig. 1, exhibiting clear difficulty-aware capabilities and reduced redundancy.

## 2 Literature Review

The advent of LRM (OpenAI et al., 2024a; Guo et al., 2025; Team, 2025) enables breakthroughs in problem solving (Sun et al., 2024b; Yu et al., 2024) with Chain-of-Thought (Wei et al., 2022; Wang et al., 2023) and test-time scaling (Snell et al., 2025). Nevertheless, this introduces the challenge of overthinking (Chen et al., 2025b; Team et al., 2025; Sui et al., 2025; Zhang et al., 2025b).

**Prompt-based methods.** Some works (Aytes et al., 2025; Ding et al., 2024; Ning et al., 2024; Han et al., 2025; Xu et al., 2025; Lee et al., 2025) instruct LRMs by various prompts to respond concisely. Specifically, Renze and Guven (2024) explicitly prompts LLMs to perform concise step-by-step reasoning. SoT (Ning et al., 2024) combines reasoning paradigms with linguistic constraints to minimize

token usage. Token-Budget (Han et al., 2025) establishes a self-estimated token budget as a constraint. Ge et al. (2025) regulates thinking token distribution by few-shot in-context learning.

**Output-based methods.** Implicit-based methods (Deng et al., 2024; Hao et al., 2024; Saunshi et al., 2025; Cheng and Durme, 2024; Shen et al., 2025a) propose to internalize reasoning steps within the latent representation, thereby shifting the cognitive process from the explicit language space to the implicit hidden space. Saunshi et al. (2025) simulates multiple steps of CoT via a looped transformer layer to deepen the reasoning process internally. Another line of work focuses on dynamic execution. Sun et al. (2024a); Xie et al. (2023); Liao et al. (2025) optimize Best-of-N (BoN) decoding with speculation and rejection. Li et al. (2024); Yang et al. (2025a); Wang et al. (2025) explores potential early exit point. NoThinking (Ma et al., 2025a) prompts reasoning models to directly output final answers without thinking.

**Model-based methods.** Fine-tuning offers a direct post-training strategy (Yang et al., 2025b). CoT-Value (Ma et al., 2025b) learns a controllable direction in the parameter space to steer generation length. C3oT (Kang et al., 2025) employs GPT-4 as a compressor while retaining key information. TokenSkip (Xia et al., 2025) fine-tunes the model to prune non-essential tokens during inference. Alternatively, reinforcement learning provides another vulnerable and expensive paradigm, often by incorporating length penalties to balance performance and efficiency (Arora and Zanette, 2025; Team et al., 2025; Aggarwal and Welleck, 2025). O1-pruner (Luo et al., 2025b) integrates a Length-Harmonizing Reward into a PPO-style loss. DAST (Shen et al., 2025b) employs SimPO (Meng and Xia, 2024) on a constructed length-preference dataset, while AdaR1 (Luo et al., 2025a) optimizes a merged hybrid long/short CoT model using DPO (Rafailov et al., 2023).

**Limitations and our approach.** Overall, existing approaches operate under a monolithic paradigm of length compression and suffer from: (1) unreliable instruction following/cognition in prompt-based methods; (2) performance and interpretability degradation in implicit output-based methods; and (3) behavioral rigidity in model-based approaches. In contrast, we reframe overthinking not just as a behavioral flaw, but as a fundamental cognitive deficit, since length is not an accurate proxy for difficulty. Our approach instills a dual-

level awareness: a global, prospective cognition of task difficulty to guide high-level strategy, and a local, retrospective in-progress perception of redundancy for fine-grained execution. This self-aware intervention mechanism cultivates the metacognitive skill of thinking about how to think efficiently, rather than simply enforcing brevity.

### 3 Preliminary and Observation

#### 3.1 Background

**LLMs.** Given a prompted question  $x$ , LLM generates a response  $y := (y_1, \dots, y_n)$  by iteratively predicting each response token  $y_i = \text{LLM}([x, y_{<i}]) \in y$  conditioned on  $x$  and generated tokens  $y_{<i} := (y_1, \dots, y_{i-1})$ . This is termed as System 1 thinking (Li et al., 2025b): [Prompt] + [Answer], where [Prompt] and [Answer] are question  $x$  and response  $y$ , which is often short and concise.

**LRMs.** Whereas in LRMs, the generation process is split into two stages, i.e., thinking and conclusion, referred to as System 2: [Prompt] +  $\langle \text{think} \rangle$ [Thoughts] $\langle / \text{think} \rangle$  + [Conclusion], where  $\langle \text{think} \rangle$  and  $\langle / \text{think} \rangle$  are two delimiters denoting the start and end of reasoning, [Thoughts] and [Conclusion] are detailed reasoning processes and concise answers, respectively. Formally, [Thoughts] reveals a trajectory of problem solving, known as CoT (Wei et al., 2022; Wang et al., 2023). At the micro level, some consecutive tokens in [Thoughts] can be aggregated into short chunk  $T_j$  based on semantics, logic and grammar:

$$[\text{Thoughts}] = T := (T_1, \dots, T_m) \quad (1)$$

Similarly, each reasoning chunk is autoregressively generated by conditioning on question  $x$  and previous chunks:  $T_i = \text{LRM}([x, T_{<i}])$ . This explicit decomposition not only improves interpretability through transparent exposure of reasoning steps but also enhances the model’s capability for complex tasks. Specifically, some reflective chunks are well structured by prefix keywords such as “Wait”, “However”, “Alternatively” (Guo et al., 2025), marking the subsequent chunks will undergo logical backtracking, self-verification, or exploration, which have been proved crucial to LRMs (Li et al., 2025a). Under this setting, we define “*overthinking*” as a phenomenon characterized by inefficiencies at both the global and local levels: globally, through an excessive number of reasoning steps ( $T_i$ ) relative to task complexity, and locally, where individual

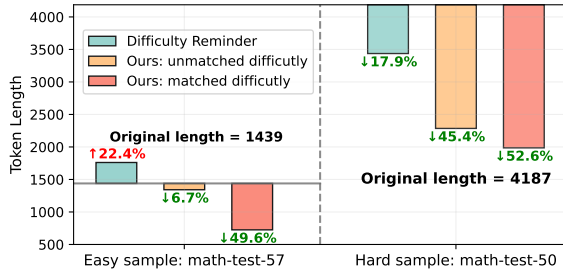


Figure 3: Internal self-hypnosis vs. external difficulty reminder for response length regulation. Our internal self-hypnosis instills a robust self-regulatory mechanism, demonstrating superior effectiveness, reliability, and instruction-following capabilities compared to conventional external prompts.

steps ( $T_i$ ) are elaborate yet logically superfluous, resulting in prominent overhead.

### 3.2 Reasoning Process Intervention

To address the overthinking issue, a straightforward approach is difficulty reminders—encouraging LRMs to self-regulate output length based on task difficulty judgment through prompts like “*Please think quickly if you find it easy*”. However, this hypothetical condition necessitates two prerequisites: accurate instruction following and robust task difficulty cognition. While prior studies (Wu et al., 2025; Han et al., 2025) have highlighted the inherent challenges of the former, we have already shown that current LRMs lack cognition of problem difficulty in Fig. 2.

Instead, we reveal that a promising solution is transferring the intervention from the prompt to the inner reasoning process. We conduct a pre-experiment in Fig. 3, pretending the model knows the difficulty level of samples. Given a simple problem, we change the model input from “*User: [Prompt]+[Question]. Assistant:*” to “*User: [Prompt]+[Question]. Assistant: This is a simple question, let’s think quickly.*”. While for hard problems, the predefined response prefix is “*It seems difficult, let’s think thoroughly.*”. Fig. 3 shows that predefined difficulty awareness largely outperforms reminder on reducing token length, especially matched awareness (“*simple, quickly*” for easy and “*difficult, thoroughly*” for hard is matched, the reverse is unmatched).

We define them as **self-hypnosis** that acts as triggers in the output, guiding the subsequent inference trajectory. Compared to explicit prompts, these built-in triggers offer three key advantages:

1) native flexibility; 2) robustness; 3) superior instruction-following capabilities. In Sec.4, we utilize the self-hypnosis mechanism to guide LRMs to generate native interventional triggers of difficulty and redundancy cognition, autonomously.

## 4 Methodology

As discussed in Sec.3.1, the challenge of overthinking stems from failures in both high-level strategy selection and fine-grained step execution. We ground our method in the dual-process theory of cognition, i.e., intuitive System 1 and analytical System 2. The strategy calls for a global, prospective intervention applied pre-generation, while the execution demands a local, retrospective mechanism for in-progress correction and refinement. Since they are complementary and distinct systems, we propose a two-stage fine-tuning framework **Think-How-to-Think (TH2T)** for mitigating overthinking, as shown in Fig. 4.

### 4.1 Stage 1: Difficulty Cognition Injection

The foundational principle of our approach is to cultivate the model’s capacity for task appraisal, which is a metacognitive prerequisite for adaptive reasoning. We define two datasets:  $\mathcal{Q}_0 = \{q_i^0 \mid i = 0, 1, 2, \dots, N\}$  for easy problems, and  $\mathcal{Q}_1 = \{q_j^1 \mid j = 0, 1, 2, \dots, M\}$  for difficult ones. Two models from the same series were selected to ensure consistency in response:  $\theta_S$  for basic LLM with prompted Short-CoT capability and  $\theta_L$  for advanced LRM. For each problem in  $\mathcal{Q}_0$ , the responses are recorded as:

$$\mathcal{P}_0^S = \{q_i^0 \in \mathcal{Q}_0 \mid \mathbb{1}[\theta_S(q_i^0)]\} \quad (2)$$

$$\mathcal{P}_0^L = \{q_i^0 \in \mathcal{Q}_0 \mid \mathbb{1}[\theta_L(q_i^0)]\} \quad (3)$$

where  $\mathbb{1}[\theta_S(q_i^0)]$  means the response of  $\theta_S$  on  $q_i^0$  is correct.  $\theta_S$  can replicate the majority correct responses of  $\theta_L$  on  $\mathcal{Q}_0$ , which means:  $\frac{|\mathcal{P}_0^S \cap \mathcal{P}_0^L|}{|\mathcal{P}_0^L|} \approx 1$ , wherein significant additional overhead yields limited performance gain. This motivates us to preserve  $\theta_S$ ’s potential for efficiency on  $\mathcal{Q}_0$  with satisfactory performance. The target reasoning pattern of  $\theta_L$  entails concise responses on  $\mathcal{Q}_0$  and detailed responses on  $\mathcal{Q}_1$ . To align with this objective, we construct a draft dataset  $\mathcal{D} = \mathcal{D}_0 \cup \mathcal{D}_1$  grounded in difficulty differentiation:

$$\mathcal{D}_0 = \{(q_i^0, \theta_S(q_i^0)) \mid q_i^0 \in \mathcal{P}_0^S\} \quad (4)$$

$$\mathcal{D}_1 = \{(q_j^1, \theta_L(q_j^1)) \mid q_j^1 \in \mathcal{P}_1^L\} \quad (5)$$

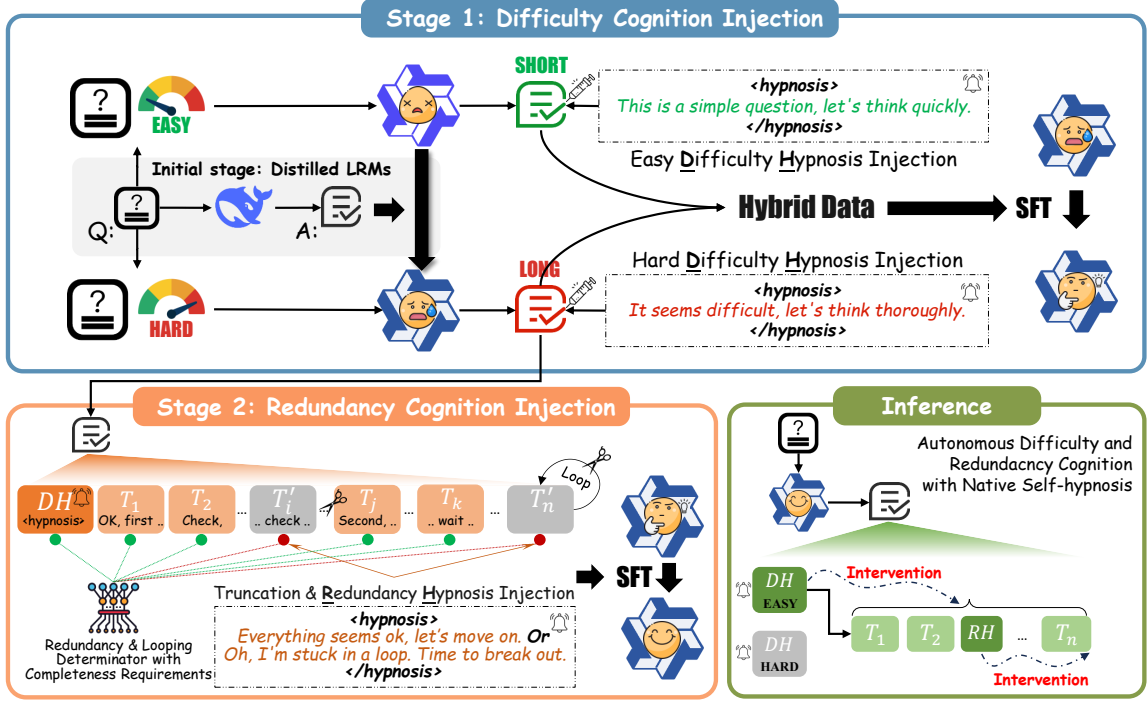


Figure 4: Framework of TH2T. **Stage 1:** fine-tune with difficulty-differentiated data with injected difficulty-hypnosis, providing global, prospective signals for strategy selection. **Stage 2:** fine-tune with truncation and redundancy-hypnosis injection, providing local, retrospective signals for in-process intervention. **Inference:** autonomous difficulty and redundancy adaptation under the intervention of native self-hypnosis.

where  $\mathcal{P}_1^L = \{q_i^1 \in \mathcal{Q}_1 \mid \mathbb{1}[\theta_L(q_j^1)]\}$ .  $\mathcal{D}_1$  aims to maintain the performance gain of  $\theta_L$  on  $\mathcal{Q}_1$ .

To enable  $\theta_L$  to perceive problem difficulty level autonomously, we introduce the trials of Sec. 3.2. We inject specific difficulty-aware cognitive Hypnosis into the prefix of answer within  $\mathcal{D}$ , which, as demonstrated before, act as an interventional signals for global, prospective strategy selection.

$$(q, \theta(q)) = (q, H + \theta(q)), \text{ for } q \in \mathcal{D} \quad (6)$$

If  $q \in \mathcal{D}_0$ , the **Difficulty Hypnosis**  $H$  is defined as follows: “`<hypnosis> This is a simple question, let's think quickly.</hypnosis>`”. If  $q \in \mathcal{D}_1$ , then “`<hypnosis> It seems difficult, let's think thoroughly.</hypnosis>`”. The markers “`<hypnosis>`” and “`</hypnosis>`” are indicators for inserted difficulty hypnosis. The injected hypnosis is recognized as an embedded difficulty label to train  $\theta_L$  to generate native difficulty cognition, stimulating a concise thinking pattern on easy tasks. The execution of stage 1 is outlined in the upper section of Fig. 4.

## 4.2 Stage 2: Redundancy Cognition Injection

We then extend difficulty-hypnosis to redundancy-hypnosis, guiding  $\theta_L$  to identify redundant structures within the reasoning process on hard tasks.

We observe severe redundant structures in  $\theta_L$ 's response on hard tasks, including superfluous reflection and check, along with a remarkably high proportion of reasoning loops in incorrect answers. According to Eq. (1), the response of  $q_j^1 \in \mathcal{D}_1$  can be aggregated into chunks based on logical semantic or grammatical rules like “`\n\n`” delimiter. The redundancy issue can be formalized as  $\theta_L(q_j^1)$ :

$$(H, T_1, \dots, T'_i, \dots, T_k, \dots, T'_n, T'_n, T'_n, \dots) \quad (7)$$

where  $H$  is the difficulty-hypnosis introduced in Stage 1,  $T_1, T_K$  are valid reasoning chunks to be kept.  $T'_i$  is the superfluous reflection structure inside the reasoning process without effective contribution to the reasoning process.  $T'_n$  is a similar reflection structure; nonetheless, it primarily manifests in the final phase of the reasoning, and the model enters an unproductive loop (trap), relentlessly iterating until  $\theta_L$  reaches the maximum output length. This practice severely degrades inference efficiency, as latency often scales super-linearly with sequence length, and the vast majority of problems do not require such exhaustive thought.

We define a **Determinator** to judge whether some  $T_m \in \theta_L(q_j^1)$  is logically equivalent to  $T'_i$  or  $T'_n$  (i.e., redundant or looping). If yes, we inject

redundancy hypnosis and truncate subsequent content. In case of  $Determinator(T_m) \equiv T'_i$ , original  $\theta_L(q_j^1)$  is refined as:

$$\theta'_L(q_j^1) = (H, T_1, T_2, \dots, T_{m-1}, \bar{H}_1) \quad (8)$$

wherein the **Redundancy Hypnosis**  $\bar{H}_1$  is defined as follows: “*\langle hypnosis \rangle Everything seems ok, let’s move on. \langle /hypnosis \rangle*”. We use GPT-4 (OpenAI et al., 2024b) to implement the determinator. In case of  $Determinator(T_m) \equiv T'_n$ , the training set lacks a sufficient number of loop structures; meanwhile, excessively deferred looping hinders the generation of shorter-length samples. Thus, we select intermediate chunk with words like “Wait”, “However”, “Alternatively” and iterate it as a simulation of looping, followed by hypnosis injection and truncation:

$$\theta'_L(q_j^1) = (H, T_1, T_2, \dots, T_m, T_m, T_m, \bar{H}_2) \quad (9)$$

wherein the **looping hypnosis**  $\bar{H}_2$  is defined as follows: “*\langle hypnosis \rangle Oh, I’m stuck in a loop. Time to break out. \langle /hypnosis \rangle*”. To determine the optimal timing for truncation, we leverage DeepSeek-R1 to justify the earliest point—a condition we term the completeness requirement—where the preceding content becomes sufficient to derive the correct answer. The choice of iterate number  $m$  can be random, and we also try a probe-like approach (Zhang et al., 2025a). The last-layer hidden states of the selected chunk tail are extracted to train a simple two-layer MLP as Determinator. Experiments indicate it yields similar results to the random approach, so we prefer the latter. Notably, we adopt truncation instead of mere substitution, as truncation allows for the construction of shorter samples. Meanwhile, generation continuity is ensured by remaining untruncated samples in the dataset, since we only apply Stage 2 to part of the samples in  $\mathcal{D}_1$ .

## 5 Experiments

### 5.1 Experiment Setup

**Models.** We do experiments on Long-CoT models, including DeepSeek-R1-Distilled-Qwen-2.5-7B/14B/32B and -Llama-8B (Guo et al., 2025). During data construction, the adopted Short-CoT model are Qwen2.5-7B/14B/32B (Qwen et al., 2025) and Llama-3.1-8B-Instruct (2024). They are  $\theta_L$  and  $\theta_S$ , respectively.

**Datasets.** For dataset construction, we mainly adopt mathematical reasoning tasks, which are

friendly for difficulty quantization. In stage 1, the GSM8K and MATH datasets are combined after filtering out errors and injecting difficulty hypnosis. In stage 2, redundancy and looping hypnosis are injected into difficult samples. The composition ratio of this heterogeneous dataset is 2:1:1, yielding  $8k$  samples. We evaluate on 4 mathematical tasks (GSM8K, MATH, AIME2024 and Omni-MATH) and 4 tasks from STEM domains, i.e., GPQA (Rein et al., 2024), LiveCodeBench (Jain et al., 2024), Olympiad (He et al., 2024), MMLU (Hendrycks et al., 2021a). Metrics include accuracy, average number of output tokens, and inference latency per sample. We also report the success rate of the difficulty awareness, redundancy, and looping optimization.

**Baselines.** We compare TH2T with the following methods: (1) **D-Prompt** (prompt-based) directly regulates response length based on self-judgment of task difficulty via the Difficulty Reminder in Prompt (in Sec.3.2). (2) **NoThinking** (Ma et al., 2025a) (output-based) bypasses the thinking process via prompting for all difficulty-level questions. (3) **Tokenskip** (Xia et al., 2025) (model-based) fine-tunes the model to prune non-essential tokens during inference, regardless of task difficulty. (4) **CoT-Valve** (Ma et al., 2025b) (model-based) learns a controllable direction in the parameter space to steer generation length based on task difficulty. (5) **AdaCtrl** (Huang et al., 2025) (model-based) uses reinforcement learning and allocates reasoning budgets based on self-assessed problem difficulty, which is more of a label-driven distillation. Further, we also compare with 5 more RL-based methods: **AdaptThink** (Zhang et al., 2025c), **DR.SAF** (Chen et al., 2025a), **DQRA** (Yan et al., 2025), **Thinkless** (Fang et al., 2025), and **Steering** (Sheng et al., 2025).

### 5.2 Results

**TH2T achieves a significant reduction in computational overhead without compromising performance.** The in-distribution performance on GSM8K and MATH is illustrated in Tab. 1. Key observations include: 1) TH2T exhibits minimal accuracy degradation; more precisely, most settings witness slight improvement. 2) TH2T achieves the best reduction in reasoning length among methods with stable accuracy—up to 74% on GSM8K and 38% on MATH with the 7B model, translating to over  $5\times$  and  $2\times$  latency speedups, respectively. 3) Greater gains are achieved on easier tasks, val-

Methods	GSM8K			MATH-500			AIME2024			OmniMath			GPQA			Average	
	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.(Gain)	Reduc.
<i>DeepSeek-R1-Distill-Qwen-7B</i>																	
Original	91.51	1057		85.2	2857		50.0	10570		45.0	5736		32.8	5349			
D-Prompt	90.59	906	-14.3%	86.0	2726	-4.6%	46.7	10209	-3.4%	45.0	5398	-5.9%	30.3	5117	-4.3%	59.7(-1.2)	-6.50%
NoThinking	88.47	243	-76.9%	79.4	699	-75.5%	40.0	4134	-60.9%	40.0	2083	-63.7%	29.3	2084	-61.0%	55.4(-5.5)	-67.6%
TokenSkip	88.01	638	-39.6%	80.2	2109	-26.2%	40.0	5559	-47.4%	38.3	3461	-39.7%	28.8	4001	-25.2%	55.1(-5.8)	-35.6%
CoT-Valve	87.90	235	-77.8%	78.6	747	-73.9%	43.3	5871	-44.5%	41.7	3311	-42.3%	28.3	3226	-39.7%	56.0(-4.9)	-55.6%
AdaCtrl	90.98	349	-	74.0	3196	-	21.3	16889	-	-	-	-	-	-	-	-	-
(w/o D.H.)	90.83	503	-47.4%	85.2	1906	-33.2%	50.1	8171	-22.7%	46.7	3571	-37.7%	33.3	4372	-18.3%	61.2(+0.3)	-31.8%
(w/o R.H.)	91.96	397	-62.4%	86.0	2265	-20.7%	48.9	8742	-17.3%	48.3	4247	-26.0%	32.8	4481	-16.2%	61.2(+0.3)	-28.6%
TH2T	92.87	275	-74.0%	86.8	1772	-38.0%	50.8	7291	-32.2%	48.3	3377	-41.1%	32.8	4166	-22.1%	62.3(+1.4)	-41.5%
<i>DeepSeek-R1-Distill-Qwen-14B</i>																	
Original	93.63	584		86.6	2306		50.0	9900		46.7	5373		50.0	3847			
D-Prompt	93.91	606	+3.9%	86.6	2297	-0.3%	50.0	10408	+5.1%	41.7	4994	-7.1%	49.0	3946	+2.6%	64.2(-1.2)	+0.8%
NoThinking	88.41	206	-64.9%	76.0	612	-73.5%	43.3	3735	-62.3%	43.3	1944	-63.8%	45.5	1454	-62.2%	59.3(-6.1)	-65.3%
TokenSkip	89.30	314	-46.2%	73.2	1356	-41.2%	46.7	6772	-31.6%	40.0	3158	-41.2%	46.9	2465	-35.9%	59.2(-6.2)	-39.2%
CoT-Valve	92.70	526	-9.9%	84.2	2111	-8.4%	40.0	7027	-29.0%	41.7	2751	-48.8%	46.5	2872	-25.3%	61.0(-4.4)	-24.3%
(w/o D.H.)	93.69	338	-42.2%	86.2	1549	-32.8%	47.4	7741	-21.8%	46.7	3745	-30.3%	50.5	2208	-42.6%	64.8(-0.6)	-33.8%
(w/o R.H.)	93.47	299	-51.2%	85.8	1706	-26.0%	50.2	8187	-17.3%	48.3	3858	-28.2%	50.5	2463	-36.0%	65.6(+0.2)	-31.8%
TH2T	94.08	246	-58.1%	86.4	1343	-41.8%	52.5	6973	-28.1%	48.3	3486	-35.1%	51.0	1901	-50.6%	66.5(+1.1)	-42.7%
<i>DeepSeek-R1-Distill-Qwen-32B</i>																	
Original	95.60	717		87.2	2357		60.0	9605		50.0	5540		52.0	4384			
D-Prompt	95.60	663	-7.6%	86.8	2198	-6.8%	56.7	9445	-1.7%	48.3	5135	-7.3%	49.5	3877	-11.6%	67.4(-1.6)	-7.03%
NoThinking	93.86	228	-68.2%	80.2	659	-72.0%	53.3	3208	-66.6%	45.0	2146	-61.3%	50.5	1627	-62.9%	64.6(-4.4)	-66.2%
TokenSkip	96.54	423	-41.0%	79.8	1567	-33.5%	50.0	5856	-39.0%	43.3	3730	-32.7%	47.4	2874	-34.4%	63.4(-5.6)	-36.1%
CoT-Valve	96.10	781	+8.9%	85.3	2263	-3.9%	53.3	7997	-16.7%	43.3	4663	-15.8%	49.0	3550	-19.0%	65.4(-3.6)	-9.30%
(w/o D.H.)	96.08	361	-49.7%	87.4	1863	-20.9%	58.6	7694	-19.9%	48.3	3432	-38.1%	51.5	2513	-42.7%	68.5(-0.5)	-34.2%
(w/o R.H.)	95.75	323	-55.0%	86.8	2114	-10.3%	60.2	8106	-15.6%	45.0	3644	-34.2%	51.0	2809	-35.9%	68.8(-0.2)	-30.1%
TH2T	97.04	263	-63.3%	88.0	1733	-26.5%	58.8	7081	-26.9%	48.3	2998	-45.9%	52.0	2367	-46.0%	69.0(+0.0)	-41.7%

Table 1: Main results of R1-Distilled 7/14/32B models on 5 benchmarks. We report accuracy, average token length and length reduction ratio. **D.H.** denotes Difficulty Hypnosis  $H$ , **R.H.** includes Redundancy Hypnosis  $\bar{H}_1$  and Looping Hypnosis  $\bar{H}_2$ . (Reduction abbreviated as “Reduc”, **Bold**: best result, underline: second-best result)

Models	GSM8K		MATH		AIME2024		Avg Gain $\Delta$		n_GPU (80G)	Training Time	Training Method
	Acc	Len	Acc	Len	Acc	Len	Acc	Len			
Original-7B	91.5	1057	85.2	2857	50.0	10570					
AdaptThink	91.0(+3.1)	309(-54.7%)	92.0(+1.8)	1875(-51.0%)	55.6(+2.1)	8599(-16.6%)	+2.3	-40.7%	8	> 10×	RL
DR.SAF	88.1(-4.3)	162(-91.2%)	88.3(-2.5)	1061(-72.5%)	50.6(+1.4)	6288(-38.3%)	-1.8	-67.3%	8	> 10×	RL
DRQA	92.7(+1.3)	1324(-23.7%)	91.4(+1.0)	3902(-23.5%)	54.7(+1.3)	10007(-27.0%)	+1.2	-24.7%	8	> 10×	RL
Thinkless	84.2(+0.7)	624(-67.4%)	81.8(-4.3)	2555(-55.0%)	27.3(-0.7)	7099(-60.7%)	-1.4	-61.0%	4	> 10×	RL
Steering	-	-	80.0(-5.2)	2050(-28.2%)	33.0(-17.0)	8900(-15.8%)	-11.2	-22.0%	8	-	-
TH2T	92.9(+1.4)	275(-74.0%)	86.8(+1.6)	1772(-38.0%)	50.8(+0.8)	7291(-32.2%)	+1.3	-48.1%	2	= 1×	SFT

Table 2: Comparisons with RL methods of 7B models, which may: 1) suffer an unacceptable drop in accuracy, 2) achieve only a limited reduction in length, or 3) incur overall > 40× training costs than ours.

idating TH2T’s ability to adapt reasoning effort based on task difficulty. 4) These observations are consistent across 7B, 14B and 32B models. As shown in Fig. 1, other methods (e.g., NoThinking and CoT-Value) achieve comparable compression ratios in token length but at the cost of severe accuracy degradation (often > 10 points), underscoring TH2T’s robustness in balanced efficiency.

**Compared with RL-based methods, TH2T achieves the best trade-off between performance gain and length reduction, while incurring the lowest training cost.** As shown in Tab.2, these methods either suffer an unacceptable drop in accuracy (e.g., -11.2% avg for Steering) or achieve a limited reduction in length (e.g., 24.7% for DRQA).

Meanwhile, these methods require 4× the GPUs and 10× the training time, yielding > 40× the total cost than ours. We believe that cost is extremely critical when making fair comparisons, especially since we are doing efficiency work.

**TH2T substantially strengthens the calibration of task difficulty cognition.** The entirety of GSM8K, MATH lv.1-2, and Omni-Math lv.1-2 are set as “Easy” while the remaining are “Hard”. As shown in Tab. 3 and Fig. 2, the original LRM reveals ambiguous recognition of task difficulty (typically < 20%), defaulting to the Medium label in the 3-option setting and lacking discriminability in the 2-option setting. In contrast, our approach yields a clear distinction with precision above 90%,

Bench	Difficulty lv.	Cognition Acc.		Confidence	
		Original	TH2T	Original	TH2T
GSM8K	all easy	14.9%	<b>98.3%</b>	15.46	<b>17.19</b>
MATH	lv.1-2 (easy)	9.8%	<b>93.4%</b>	16.09	<b>16.97</b>
	lv.4-5 (hard)	12.0%	<b>91.5%</b>	15.72	<b>16.25</b>
AIME2024	all hard	16.7%	<b>96.7%</b>	16.64	<b>16.89</b>
OmniMath	lv.1-2 (easy)	20.0%	<b>96.7%</b>	16.41	<b>17.05</b>
	lv.4-5 (hard)	13.3%	<b>93.3%</b>	16.94	<b>17.76</b>

Table 3: Difficulty Cognition. We compare the correctness of difficulty assessment by R1-Distilled-7B and our TH2T-7B, further examining the first-token confidence as an internal signal for self-consistent metacognition.

Bench	Split	Reflective		Looping	
		Original	TH2T	Original	TH2T
GSM8K	Overall	10.3	<b>1.2</b>	2.0%	<b>0.0%</b>
	-correct	4.8	<b>0.0</b>	0.4%	<b>0.0%</b>
	-incorrect	69.4	<b>10.8</b>	22.3%	<b>0.0%</b>
MATH	Overall	27.0	<b>12.3</b>	10.4%	<b>4.0%</b>
	-correct	14.9	<b>6.8</b>	1.9%	<b>1.0%</b>
	-incorrect	96.5	<b>39.1</b>	59.5%	<b>21.4%</b>

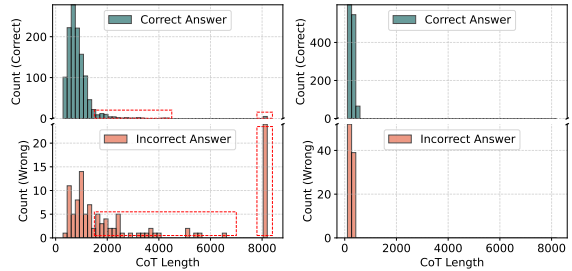
Table 4: Statistics of reflective and looping structure counts (Original R1-Distilled-Qwen-7B vs TH2T-7B). Reflective structures include chunks containing “Wait”, “Alternatively”, or “But”. Looping structures are defined as extremely long answers ending with repetitive patterns up to the maximum generation length. Tasks are further divided based on correctness.

evidencing effective calibration of task difficulty cognition enabled by the first-stage training.

**TH2T substantially enhances cognition of redundancy.** Tab. 4 reports the average number of reflective and looping structures, comparing R1-Distilled-7B and TH2T-7B. We observe that: 1) TH2T effectively suppresses the number of reflection chunks by 10 $\times$  on easy tasks and 3 $\times$  on hard tasks, suggesting that extensive reflection is superfluous for strong reasoning. 2) TH2T substantially alleviates the looping problem. Notably, 59.5% of incorrect answers by the original model on MATH contain tail recursion, which is reduced to 21.4% under TH2T. As shown in Fig. 5(b), we observe zero loop structure in GSM8K. This distinct pattern of short-length concentration highlights the efficacy of Stage-2 of TH2T, yielding more concise and efficient reasoning structures.

### 5.3 Analysis

**Generalization Ability.** To validate the generalization on OOD tasks, we test on 2 more mathematical benchmarks in Tab. 1, i.e., AIME2024



(a) Original-7B

(b) TH2T-7B

Figure 5: Statistics of length distribution on GSM8K. Our approach eliminates longer responses, especially repetitive ones that reach maximum generation length, presenting more efficiency in terms of redundancy.

and OmniMath. To validate the generalization on tasks beyond math, we report results on 4 tasks from STEM domain, including QA, coding, science, and multi-tasks. As shown in Tab. 6, consistent with the main results, our approach preserves accuracy while yielding substantial token reduction over 30%.

#### Internal states validate the enhanced cognition.

To examine whether our method instills a genuine sense of difficulty cognition, we investigate the model’s internal state in Tab.3. We posit that the confidence (Fu et al., 2025) of the first token acts as an intrinsic, metacognitive signal of the model’s perceived task difficulty. Our method exhibits higher confidence, aligning with a concurrent rise in difficulty assessment accuracy. This consistency serves as a clear, self-consistent internal signal, providing strong validation that the enhanced cognition is explicitly represented in internal state.

#### Ablation on two-stage design.

We explore two more settings: 1) without Difficulty Hypnosis in Stage-1, and 2) without Redundancy Hypnosis in Stage-2, which are annotated as *w/o D.H.* and *w/o R.H.* in Tab. 1. We observe that D.H. introduces more token reduction on easy task GSM8K than on hard task MATH (e.g., 26.6% vs 4.8%, on 7B model), while R.H. exerts a greater impact on MATH (11.6% vs 17.3%). As we illustrate in Sec.4, we design Difficulty Cognition Injection Stage to build native perception on easy tasks and the Redundancy Cognition Injection Stage to mitigate non-essential structures that frequently occur in hard tasks. This phenomenon also holds for the other three datasets, demonstrating the rationality of our two-stage framework.

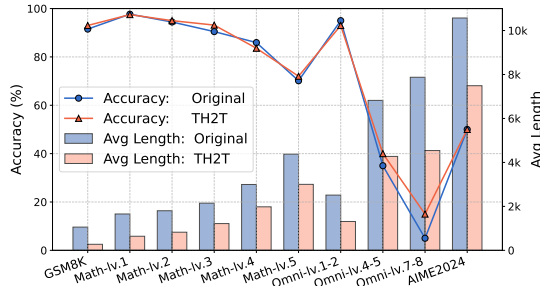


Figure 6: Breakdown of the acc and average length distribution across the spectrum of difficulty levels.

**Acc and length across difficulty levels.** We report a breakdown of the accuracy and average token length distribution across the complete spectrum of difficulty levels on 4 tasks in Fig.6. We compare the original R1-Distilled-Qwen-7B with TH2T-7B, where higher numerical levels indicate greater question difficulty. Our observations are as follows: 1) Both models produce progressively longer and more elaborate responses with increasing difficulty. 2) Across all difficulty levels and tasks, TH2T consistently reduces response length while maintaining stable performance. 3) Length reduction is most pronounced for easier questions.

## 6 Conclusion

This paper presents TH2T, a novel two-stage fine-tuning strategy for large reasoning models (LRMs) that effectively addresses the challenge of overthinking. Our empirical analysis indicates that current LRMs are largely limited to recognizing task difficulty, exhibiting a necessity for aligning with human *System 1* and *System 2* thinking patterns. Motivated by this, TH2T progressively enables native difficulty and redundancy cognition based on a hybrid dataset and proposed self-hypnosis mechanism. To the best of our knowledge, we are the first to inject difficulty/redundancy hypnosis as built-in triggers for global, prospective, or local, retrospective interventional signals. Experiments demonstrate that TH2T significantly reduces inference length without compromising performance, exhibiting clear sensitivity to difficulty and redundancy.

## Limitations

First, our current approach primarily focuses on standard fine-tuning using the generated hybrid dataset. While this method is effective within our current framework, it may not fully exploit the potential for optimized efficiency. Reinforcement

learning presents a promising avenue for further performance gains, although it can be unstable and extremely computationally demanding. Second, the scope of this study is confined to language-only models. Consequently, the efficacy of our approach in multimodal architectures remains unexplored. Investigating whether our method can be adapted to achieve competitive performance on vision-language tasks constitutes an interesting direction for further research.

## Acknowledgement

This research was supported by fundings from the Hong Kong RGC General Research Fund (152228/23E, 162161/24E, 162116/25E, 162180/25E), National Natural Science Foundation of China (NSFC) Key Program (No.62532005), Collaborative Research Fund (No. C1042-23GF, No. C5097-25G), NSFC/RGC Collaborative Research Scheme (Grant No. 62461160332 & CRS HKUST602/24), Research Impact Fund (No. R5011-23F), Areas of Excellence Scheme (AoE/E-601/22-R), and the InnoHK (HKGAI).

## References

- Pranjal Aggarwal and Sean Welleck. 2025. [L1: Controlling how long a reasoning model thinks with reinforcement learning](#). *Preprint*, arXiv:2503.04697.
- Daman Arora and Andrea Zanette. 2025. [Training language models to reason efficiently](#). *Preprint*, arXiv:2502.04463.
- Simon A. Aytes, Jinheon Baek, and Sung Ju Hwang. 2025. [Sketch-of-thought: Efficient llm reasoning with adaptive cognitive-inspired sketching](#). *Preprint*, arXiv:2503.05179.
- Qiguang Chen, Dengyun Peng, Jinhao Liu, HuiKang Su, Jiannan Guan, Libo Qin, and Wanxiang Che. 2025a. [Aware first, think less: Dynamic boundary self-awareness drives extreme reasoning efficiency in large language models](#). *arXiv preprint arXiv:2508.11582*.
- Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu, Mengfei Zhou, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025b. [Do not think that much for 2+3=? on the overthinking of o1-like llms](#). *Preprint*, arXiv:2412.21187.
- Jeffrey Cheng and Benjamin Van Durme. 2024. [Compressed chain of thought: Efficient reasoning through dense representations](#). *Preprint*, arXiv:2412.13171.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias

- Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Yuntian Deng, Yejin Choi, and Stuart Shieber. 2024. [From explicit cot to implicit cot: Learning to internalize cot step by step](#). *Preprint*, arXiv:2405.14838.
- Mengru Ding, Hanmeng Liu, Zhizhang Fu, Jian Song, Wenbo Xie, and Yue Zhang. 2024. [Break the chain: Large language models can be shortcut reasoners](#). *Preprint*, arXiv:2406.06580.
- Anastasia Efklides. 2011. Interactions of metacognition with motivation and affect in self-regulated learning: The masrl model. *Educational psychologist*, 46(1):6–25.
- Jonathan St B T Evans and Keith E Stanovich. 2013. Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3):223–241.
- Gongfan Fang, Xinyin Ma, and Xinchao Wang. 2025. [Thinkless: LLM learns when to think](#). In *The Thirtieth Annual Conference on Neural Information Processing Systems*.
- John H Flavell. 1979. Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry. *American Psychologist*, 34(10):906–911.
- Yichao Fu, Xuwei Wang, Yuandong Tian, and Jiawei Zhao. 2025. [Deep think with confidence](#). *arXiv preprint arXiv:2508.15260*.
- Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Runxin Xu, Zhengyang Tang, Benyou Wang, Daoguang Zan, Shanghaoran Quan, Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. 2024. [Omni-math: A universal olympiad level mathematic benchmark for large language models](#). *Preprint*, arXiv:2410.07985.
- Yuyao Ge, Shenghua Liu, Yiwei Wang, Lingrui Mei, Lizhe Chen, Baolong Bi, and Xueqi Cheng. 2025. [Innate reasoning is not enough: In-context learning enhances reasoning large language models with less overthinking](#). *Preprint*, arXiv:2503.19602.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and et al. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*.
- Tingxu Han, Zhenting Wang, Chunrong Fang, Shiyu Zhao, Shiqing Ma, and Zhenyu Chen. 2025. [Token-budget-aware LLM reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24842–24855, Vienna, Austria. Association for Computational Linguistics.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2024. [Training large language models to reason in a continuous latent space](#). *Preprint*, arXiv:2412.06769.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. [Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021a. [Measuring massive multitask language understanding](#). *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, and Dawn Song. 2021b. [Measuring mathematical problem solving with the math dataset](#). *NeurIPS*.
- Shijue Huang, Hongru Wang, Wanjuan Zhong, Zhaochen Su, Jiazhan Feng, Bowen Cao, and Yi R Fung. 2025. [Adactrl: Towards adaptive and controllable reasoning via difficulty-aware budgeting](#). *arXiv preprint arXiv:2505.18822*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. [Livecodebench: Holistic and contamination free evaluation of large language models for code](#). *arXiv preprint arXiv:2403.07974*.
- Daniel Kahneman. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Yu Kang, Xianghui Sun, Liangyu Chen, and Wei Zou. 2025. [C3ot: Generating shorter chain-of-thought without compromising effectiveness](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 24312–24320.
- Ayeong Lee, Ethan Che, and Tianyi Peng. 2025. [How well do llms compress their own chain-of-thought? a token complexity approach](#). *Preprint*, arXiv:2503.01141.
- Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xi-angxi Mo, Eric Tang, Sumanth Hegde, Kourosh Hakhmaneshi, Shishir G. Patil, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. 2025a. [Llms can easily learn to reason from demonstrations structure, not content, is what matters!](#) *Preprint*, arXiv:2502.07374.

- Yiwei Li, Peiwen Yuan, Shaoxiong Feng, Boyuan Pan, Xinglin Wang, Bin Sun, Heda Wang, and Kan Li. 2024. [Escape sky-high cost: Early-stopping self-consistency for multi-step reasoning](#). In *The Twelfth International Conference on Learning Representations*.
- Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhiwei Li, Bao-Long Bi, Ling-Rui Mei, Junfeng Fang, Xiao Liang, Zhi-jiang Guo, and 2 others. 2025b. [From system 1 to system 2: A survey of reasoning large language models](#). *Preprint*, arXiv:2502.17419.
- Baohao Liao, Yuhui Xu, Hanze Dong, Junnan Li, Christof Monz, Silvio Savarese, Doyen Sahoo, and Caiming Xiong. 2025. [Reward-guided speculative decoding for efficient LLM reasoning](#). In *Forty-second International Conference on Machine Learning*.
- Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. 2025. [Understanding r1-zero-like training: A critical perspective](#). *Preprint*, arXiv:2503.20783.
- Haotian Luo, Haiying He, Yibo Wang, Jinluan Yang, Rui Liu, Naiqiang Tan, Xiaochun Cao, Dacheng Tao, and Li Shen. 2025a. [Ada-r1: Hybrid-cot via bi-level adaptive reasoning optimization](#). *Preprint*, arXiv:2504.21659.
- Haotian Luo, Li Shen, Haiying He, Yibo Wang, Shiwei Liu, Wei Li, Naiqiang Tan, Xiaochun Cao, and Dacheng Tao. 2025b. [O1-pruner: Length-harmonizing fine-tuning for o1-like reasoning pruning](#). *Preprint*, arXiv:2501.12570.
- Wenjia Ma, Jingxuan He, Charlie Snell, Tyler Griggs, Sewon Min, and Matei Zaharia. 2025a. [Reasoning models can be effective without thinking](#). *Preprint*, arXiv:2504.09858.
- Xinyin Ma, Guangnian Wan, Runpeng Yu, Gongfan Fang, and Xinchao Wang. 2025b. [CoT-valve: Length-compressible chain-of-thought tuning](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6025–6035, Vienna, Austria. Association for Computational Linguistics.
- Yu Meng and Mengzhou Xia. 2024. [Simpo: Simple preference optimization with a reference-free reward](#). *Advances in Neural Information Processing Systems*, 37:124198–124235.
- Xuefei Ning, Zinan Lin, Zixuan Zhou, Zifu Wang, Huazhong Yang, and Yu Wang. 2024. [Skeleton-of-thought: Prompting LLMs for efficient parallel generation](#). In *The Twelfth International Conference on Learning Representations*.
- OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, and et al. 2024a. [Openai o1 system card](#). *Preprint*, arXiv:2412.16720.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, and et al. 2024b. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, and et al. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). *Advances in Neural Information Processing Systems*, 36:53728–53741.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. [Gpqa: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling*.
- Matthew Renze and Erhan Guven. 2024. [The benefits of a concise chain of thought on problem-solving in large language models](#). In *2024 2nd International Conference on Foundation and Large Language Models (FLLM)*, pages 476–483. IEEE.
- Nikunj Saunshi, Nishanth Dikkala, Zhiyuan Li, Sanjiv Kumar, and Sashank J. Reddi. 2025. [Reasoning with latent thoughts: On the power of looped transformers](#). In *The Thirteenth International Conference on Learning Representations*.
- Xuan Shen, Yizhou Wang, Xiangxi Shi, Yanzhi Wang, Pu Zhao, and Jiuxiang Gu. 2025a. [Efficient reasoning with hidden thinking](#). *Preprint*, arXiv:2501.19201.
- Yi Shen, Jian Zhang, Jieyun Huang, Shuming Shi, Wenjing Zhang, Jiangze Yan, Ning Wang, Kai Wang, Zhaoxiang Liu, and Shiguo Lian. 2025b. [Dast: Difficulty-adaptive slow-thinking for large reasoning models](#). *Preprint*, arXiv:2503.04472.
- Leheng Sheng, An Zhang, Zijian Wu, Weixiang Zhao, Changshuo Shen, Yi Zhang, Xiang Wang, and Tat-Seng Chua. 2025. [On reasoning strength planning in large reasoning models](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Charlie Victor Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2025. [Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning](#). In *The Thirteenth International Conference on Learning Representations*.
- Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Hanjie Chen, and Xia

- Hu. 2025. [Stop overthinking: A survey on efficient reasoning for large language models](#). *Preprint*, arXiv:2503.16419.
- Hanshi Sun, Momin Haider, Ruiqi Zhang, Huitao Yang, Jiahao Qiu, Ming Yin, Mengdi Wang, Peter Bartlett, and Andrea Zanette. 2024a. [Fast best-of-n decoding via speculative rejection](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Jiankai Sun, Chuanyang Zheng, Enze Xie, Zhengying Liu, Ruihang Chu, Jianing Qiu, Jiaqi Xu, Mingyu Ding, Hongyang Li, Mengzhe Geng, and et al. 2024b. [A survey of reasoning with foundation models](#). *Preprint*, arXiv:2312.11562.
- Kimi Team, Angang Du, Bofei Gao, Bawei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, and 1 others. 2025. [Kimi k1.5: Scaling reinforcement learning with llms](#). *Preprint*, arXiv:2501.12599.
- Qwen Team. 2025. Qwq-32b: Embracing the power of reinforcement learning. <https://qwenlm.github.io/blog/qwq-32b/>.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Yiming Wang, Pei Zhang, Siyuan Huang, Baosong Yang, Zhuosheng Zhang, Fei Huang, and Rui Wang. 2025. [Sampling-efficient test-time scaling: Self-estimating the best-of-n sampling in early decoding](#). *Preprint*, arXiv:2503.01422.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in neural information processing systems*, 35:24824–24837.
- Tong Wu, Chong Xiang, Jiachen T. Wang, G. Edward Suh, and Prateek Mittal. 2025. [Effectively controlling reasoning models through thinking intervention](#). *Preprint*, arXiv:2503.24370.
- Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. 2025. [Tokenskip: Controllable chain-of-thought compression in llms](#). *Preprint*, arXiv:2502.12067.
- Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. 2025. [Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning](#). *Preprint*, arXiv:2502.14768.
- Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, and Michael Xie. 2023. [Self-evaluation guided beam search for reasoning](#). *Advances in Neural Information Processing Systems*, 36:41618–41650.
- Silei Xu, Wenhao Xie, Lingxiao Zhao, and Pengcheng He. 2025. [Chain of draft: Thinking faster by writing less](#). *Preprint*, arXiv:2502.18600.
- Kaiwen Yan, Xuanqing Shi, Hongcheng Guo, Wenxuan Wang, Zhuosheng Zhang, and Chengwei Qin. 2025. [Drqa: Dynamic reasoning quota allocation for controlling overthinking in reasoning large language models](#). *arXiv preprint arXiv:2508.17803*.
- Chenxu Yang, Qingyi Si, Yongjie Duan, Zheliang Zhu, Chenyu Zhu, Qiaowei Li, Zheng Lin, Li Cao, and Weiping Wang. 2025a. [Dynamic early exit in reasoning models](#). *Preprint*, arXiv:2504.15895.
- Wenkai Yang, Shuming Ma, Yankai Lin, and Furu Wei. 2025b. [Towards thinking-optimal scaling of test-time compute for llm reasoning](#). *Preprint*, arXiv:2502.18080.
- Fei Yu, Hongbo Zhang, Prayag Tiwari, and Benyou Wang. 2024. [Natural language reasoning, a survey](#). *ACM Computing Surveys*, 56(12):1–39.
- Anqi Zhang, Yulin Chen, Jane Pan, Chen Zhao, Aurojit Panda, Jinyang Li, and He He. 2025a. [Reasoning models know when they're right: Probing hidden states for self-verification](#). *Preprint*, arXiv:2504.05419.
- Haoyue Zhang, Hualei Zhang, Xiaosong Ma, Jie Zhang, and Song Guo. 2025b. [Lazyeviction: Lagged kv eviction with attention pattern observation for efficient long reasoning](#). *arXiv preprint arXiv:2506.15969*.
- Jiajie Zhang, Nianyi Lin, Lei Hou, Ling Feng, and Juanzi Li. 2025c. [Adaptthink: Reasoning models can learn when to think](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 3716–3730.

# Appendices

## A More Experimental Results

### A.1 Results on Llama Model

We provide detailed experimental results of the R1-Distilled-Llama-8B model evaluated on 4 mathematical benchmarks (i.e., GSM8K, MATH, AIME2024, and OmniMath) and 1 QA benchmark (i.e., GPQA). We report accuracy and average token length for comparison in Tab. 5.

### A.2 Ablation on Training Steps.

We examine the dynamics of accuracy and average token length during the training process (in Fig. 7), and obtain two interesting findings: 1) As training steps increase, accuracy initially declines, then returns to the original level, with occasional positive gains during the process. 2) For hard tasks like MATH, the response length exhibits a reverse increase after exceeding a certain number of steps. This inspires us to set the training steps to 6400 to prevent model degradation.

### A.3 Visualization of the Effect on Token Length

We collect the first 50 samples from GSM8K and MATH test set in sequential order, and compare the output token length distribution between our TH2T and the original DeepSeek-R1-Distill-Qwen-7B in Fig. 8. We can see TH2T significantly reduces response length.

## B More Experimental Details

### B.1 Dataset Construction

The datasets used in this paper incorporate difficulty level labeling as a key consideration, so we mainly use mathematical reasoning tasks. Our TH2T is a two-stage strategy, with tailored dataset construction for each phase, followed by sequential fine-tuning. 1) During the first stage, we use the GSM8K and MATH training datasets to obtain  $\mathcal{D}_0$  and  $\mathcal{D}_1$ , respectively. The former is generally considered simple, whereas the latter is regarded as difficult and is explicitly categorized into difficulty levels ranging from 1 to 5, and we pick samples between lv.3 and lv.5. After filtering out errors,  $\mathcal{D}_0$  and  $\mathcal{D}_1$  are balanced at a 1:1 ratio. After injecting difficulty hypnosis  $H$ , we construct  $\mathcal{D}$  of 6.4K samples. 2) In Stage-2,  $\mathcal{D}_1$  are injected with

redundancy hypnosis  $\bar{H}_1$  and looping hypnosis  $\bar{H}_2$ , forming a final training dataset. We observe that the meticulous composition ratio of the constructed dataset does not significantly affect the final performance. Hence, we combine  $\mathcal{D}_0$  and  $\mathcal{D}_1$  in a ratio of 1:1. Meanwhile, the redundancy hypnosis and looping hypnosis are also injected into  $\mathcal{D}_1$  in a ratio of 1:1. The overall composition is set as 2:1:1, yielding  $8k$  samples in total.

During data construction, we clarify our two designs to prevent a potential lack of logical coherence caused by truncation rules.

- **Completeness requirement.** As described in Section 4.2, lexical keywords are just candidate indicators. The strict "Completeness Requirement" ensures the Determinator only truncates when reasoning is logically sufficient, discarding broken logical chains.
- **Data Composition.** Intervened samples are the minority (ratio 2:1:1). The majority of full-length CoT ensures global coherence. Maintained acc (Tab.1) empirically proves logic remains intact.

### B.2 Metrics

Main results are assessed on accuracy and the average number of output tokens. To substantiate the research objectives, we incorporate the success rate of the difficulty awareness. We also report the redundancy cognition optimization, which includes the reduction of redundant structures (e.g., "Wait", "Alternatively", or "Check"), as well as the reflection looping ratio. To examine whether our method instills a genuine sense of difficulty cognition, we investigate the model's internal state, i.e., the confidence value of the first generated token.

### B.3 Implementation Details

We adopt LoRA, a widespread LLM fine-tuning approach, to train our models with a rank  $r = 8$ , scaling parameter  $\alpha = 16$ , and learning rate  $r = 1 \times 10^{-5}$ . All experiments are conducted on 4xNVIDIA H800 GPU (80GB). The maximum generation length is set to 8192 for GSM8K and MATH500, and 16,384 for AIME2024, OmniMath, and GPQA. For the decoding strategy, we employ greedy decoding. The random seed is fixed to 42.

Methods	GSM8K		MATH-500		AIME2024		OmniMath		GPQA		Average	
	Acc	Len	Acc	Len	Acc	Len	Acc	Len	Acc	Len	Acc	Len
<i>DeepSeek-R1-Distill-Llama-8B</i>												
Original	75.36	260	61.8	701	3.3	1378	26.67	844	34.34	704	40.29	777
D-Prompt	75.14	261	<b>62.2</b>	596	6.6	1100	28.33	808	34.34	669	41.32	687
NoThinking	75.24	257	<b>62.2</b>	580	6.6	1019	28.33	714	32.32	619	40.94	638
<b>TH2T</b>	<b>75.58</b>	<b>245</b>	61.6	<b>498</b>	<b>13.3</b>	<b>861</b>	<b>40.00</b>	<b>561</b>	<b>40.91</b>	<b>526</b>	<b>46.28</b>	<b>538</b>

Table 5: Results of R1-Distilled-Llama-8B model on 4 mathematical benchmarks (i.e., GSM8K, MATH, AIME2024, and OmniMath) and 1 QA benchmark (i.e., GPQA). We report accuracy, average token length for comparison.

Methods	GPQA			LiveBenchCode			Olympiad			MMLU			Average		
	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.	Acc.	Len.	Reduc.
Original-7B	32.8	5349		43.5	5890		50.7	5213		63.4	4132		47.6	5146	
D-Prompt	30.3	5117	-4.30%	42.3	5472	-7.10%	49.4	4676	-10.3%	62.5	3802	-6.70%	46.1	4766	-7.10%
NoThinking	29.3	2084	-61.0%	37.2	1991	-66.2%	43.4	1491	-71.4%	52.2	1319	-66.9%	40.5	1721	-66.4%
<b>TH2T</b>	32.8	4166	-22.1%	44.0	3852	-34.6%	52.0	3122	-40.1%	64.2	2588	-37.9%	48.3	3432	-33.7%

Table 6: Generalization on tasks beyond math. We report results on 4 benchmarks across QA (GPQA), coding (LiveCodeBench), science (Olympiad), and multi-task STEM (MMLU) tasks.

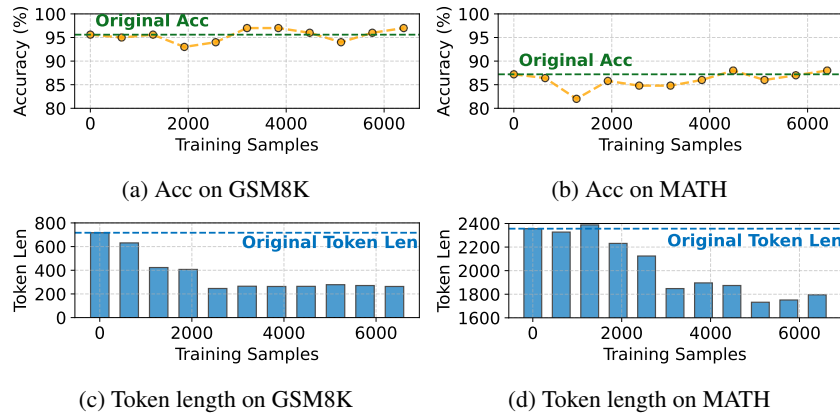


Figure 7: Dynamics of accuracy and average token length during the training process (R1-Distill-Qwen-32B). We find: 1) As training steps increase, accuracy initially declines, then returns to the original level, with occasional positive gains during the process. 2) For hard tasks like MATH, the response length exhibits a reverse increase after exceeding a certain number of steps.

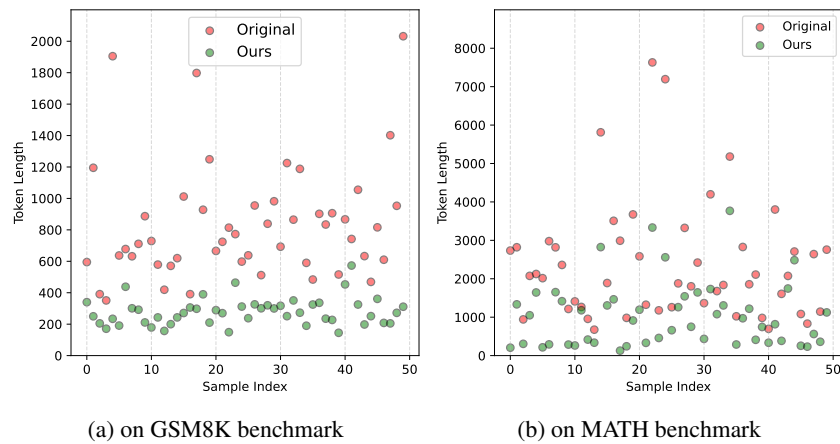


Figure 8: Visualization of the effect on token length. We can see TH2T significantly reduces response length.

Except for AIME2024, as evaluating on the 30-question AIME dataset with greedy decoding leads to high variance, we instead report the average accuracy across eight sampled paths.

## C More Explanation on Motivation

We investigate the efficiency issue from a suspicion: *Are current LRMs treating various task types equally with identical response strategy?* We compare the accuracy gain and inference length scaling ratio between Short-CoT model (Qwen2.5-7B) and Long-CoT model (DeepSeek-R1-Distill-Qwen2.5-7B) on easy GSM8K and hard MATH-500 benchmarks. From Fig. 9, we obtain **three findings**:

- Across the easy and difficult tasks, the current reasoning model shows similar token length scaling ratio (i.e., both around 400%), which verifies our initial suspicion that LRMs are allocating similar additional computing resources to tasks with varying difficulty levels.
- Despite incurring minimal costs, basic LLM can still replicate the majority of correct responses (89.9%) of LRM’s on the easy tasks, which is why we claimed that  $\theta_S$  can replicate the majority correct responses of  $\theta_L$  on  $\mathcal{Q}_0$  in Sec.4.1.
- Compared with difficult tasks, the easy task only witnesses a minor accuracy gain (5.3% vs 22.4%) under the same computation resource allocation strategy, which is not what we expected. We aim to maintain this marginal performance improvement while keeping the associated token length increase at a minimal level.

These raise a reasonable speculation: **LRMs fall short of perceiving task complexity, since massive inference resources on easy ones are superfluous and lead to limited performance gain.** As discussed in Results on Redundancy Cognition of the main paper, we observe severe redundant structures in LRM’s responses to hard tasks, including superfluous reflection along with a remarkably high proportion of reasoning loops in incorrect answers. These findings lead to the **two objectives** of this paper:

- Enabling LRMs to think how to think, i.e., autonomously recognizing and adapting to variant levels of problem difficulty (results and

analysis are presented in the Results of Difficulty Cognition section of the main paper).

- Mitigating LRM overthinking by promoting redundancy cognition of internal reasoning structure (results and analysis are presented in the Results of Redundancy Cognition section of the main paper).

## D Prompts

For reproducibility, we provide the complete list of prompts adopted in our experiments.

### D.1 Difficulty level evaluation

Fig. 10 shows the prompt template we used during difficulty level evaluation.

### D.2 Baseline

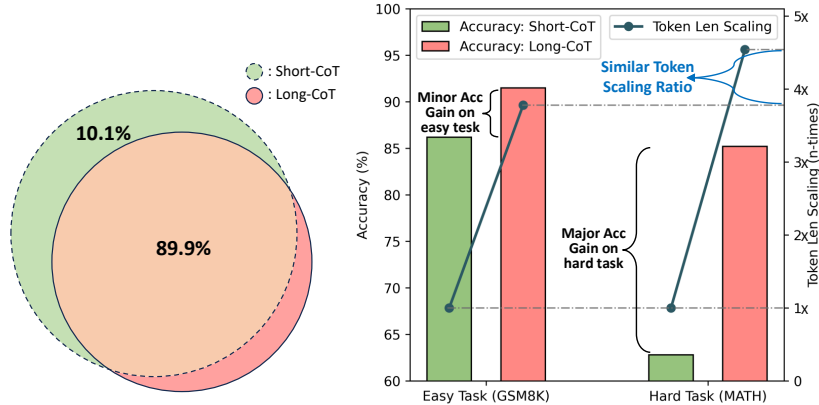
Fig. 11 is the difficulty reminder for the baseline **D-Prompt**, which is also adopted in the pre-experiment. Fig. 12 is the setting of baseline **No-Thinking**.

## E Design of Determinator

During Stage 2 of TH2T, we design a Determinator. To identify redundancy point, i.e.,  $Determinator(T_m) \equiv T'_i$ , the contribution effectiveness of reflective structures is assessed as shown in Fig. 14 and Fig. 15. Furthermore, when injecting hypnosis, the completeness requirement in Fig. 13 is adopted.

## F Case Study

Three concrete cases are provide in Fig. 16, Fig. 17 and 18.



(a) Overlap of correct responses on GSM8K (b) Accuracy gain and token length scaling

Figure 9: Short-CoT model (Qwen2.5-7B) vs LRM (DeepSeek-R1-Distill-Qwen2.5-7B). Left: Short-CoT replicates up to 89.9% correct responses of Long-CoT on the easy task. Right: under a similar high token length scaling ratio (around 400%), the easy task only witnesses a minor accuracy gain compared with the hard task. This comparison raises a reasonable speculation that LRMs fall short of perceiving task complexity, since massive inference resources on easy ones are superfluous and lead to limited performance gain.

**Prompt Template for Difficulty Level Cognition**

```

<|im_start|>system\nYou are a helpful assistant.<|im_end|>\n<|im_start|>user\nHere is a question.\n" + f"{mess['content']}" + ". Please evaluate the difficulty level of the above question on a scale from 1 to 10. A score of 1 represents very simple questions, a score of 10 represents very difficult questions, while a score of 5 represents medium level. Consider the complexity, required knowledge, and effort needed to solve the question based on your cognition when assigning a score. Do not answer the given question. Instead, just give your evaluation of difficulty level without any explanation. Put your final answer within \boxed{ }. <|im_end|>\n<|im_start|>assistant\n<think>\nOkay, I have finished evaluation. \n</think>\n\nThe difficulty level of the above question is"

```

Figure 10: Prompt Template for Difficulty Level Evaluation

**Baseline: Difficulty Reminder**

```

<|im_start|>system\nYou are a helpful assistant.<|im_end|>\n<|im_start|>user\nPlease reason step by step, and put your final answer within \boxed{ }. Please think quickly if you find it easy. Otherwise, take your time.\n" + f"{mess['content']}" <|im_end|>\n<|im_start|>assistant\n

```

Figure 11: Setting of baseline – Difficulty Reminder

**Baseline: NoThinking**

```

<|im_start|>system\nYou are a helpful assistant.<|im_end|>\n<|im_start|>user\nPlease reason step by step, and put your final answer within \boxed{ }.\n" + f"{mess['content']}" <|im_end|>\n<|im_start|>assistant\n<think>\nOkay, I have finished thinking.</think>\n\n

```

Figure 12: Setting of baseline – NoThinking

### Determinator: Completeness Requirements

Given several chunks of a reasoning trace, along with a ground-truth answer, independently evaluate each chunk. If a chunk reaches a result at the end, return the intermediate result; otherwise, return None if the chunk does not contain an intermediate result (e.g., pure reflections).

Then, if an intermediate answer exists, compare it to the ground-truth answer. If the intermediate result in the chunk equals the ground-truth answer, return True; if the intermediate result in the chunk does not equal the ground-truth answer, return False; if no intermediate answer exists, return None. Output in JSON format:

```
[  
  {"id": "1", "result": "6 + 9i" / None, "correctness": True / False / None},  
  ...  
]
```

**Input chunks:** {reasoning trace}

**Ground-truth answer:** {answer}

Figure 13: Design of Determinator – Completeness Requirements.

### Determinator: Redundant Reflection

You are tasked with refining a step-by-step mathematical reasoning process by identifying and removing unnecessary backtracking or exploration that does not contribute to the final solution. These sections often begin with phrases such as "wait", "alternatively", "Wait, maybe I need", "But wait, the problem says", "Wait, but maybe there's another way", "Alternatively, maybe", "Alternatively, we can", or similar expressions. These backtracking or exploration attempts are redundant if the reasoning following them confirms the prior steps were already correct or does not introduce meaningful progress. Your task is to:

- Detect the redundant backtracking or exploration sections.
- Replace these sections with the phrase: `</hypnosis>Everything seems ok, let's move on.</hypnosis>`
- Ensure the remaining text is coherent and preserves the logical flow.
- Keep all other parts of the reasoning process unchanged.
- If the redundant section spans multiple sentences, ensure all related sentences are replaced with the above phrase for semantic consistency.

Here is an example for removing backtracking :

#### Input Example:

Therefore, there are 2 distinct ordered pairs.

But wait, hold on. The problem states "distinct ordered pairs". The pairs (5,20) and (6,12) are indeed all different when considering order. Let's test each pair to ensure they satisfy  $1/m + 1/n = 1/4$ .

1. (5,20):  $1/5 + 1/20 = 4/20 + 1/20 = 5/20 = 1/4$ . Correct.

2. (6,12):  $1/6 + 1/12 = 2/12 + 1/12 = 3/12 = 1/4$ . Correct.

All pairs check out. Therefore, the answer should be 5.

#### Output Example:

Therefore, there are 2 distinct ordered pairs.

`</hypnosis>Everything seems ok, let's move on.</hypnosis>`

All pairs check out. Therefore, the answer should be 5.

Here is an example for removing exploration :

#### Input Example:

Yep, that matches the product given. So the sum is indeed 45.  
Alternatively, maybe there's a way without solving the quadratic. Since 506 is between  $20^2=400$  and  $25^2=625$ . Let's see,  $20*21=420$ ,  $21*22=462$ ,  $22*23=506$ . Oh, that's it! So the numbers are 22 and 23. Their sum is 45. So either way, same answer. Next, we going to continue calculation.

#### Output Example:

Yep, that matches the product given. So the sum is indeed 45.  
`</hypnosis>Everything seems ok, let's move on.</hypnosis>` Next, we going to continue calculation.

Now, apply this transformation to the provided text. Only detect and replace backtracking or exploration sections. Do not modify any other parts of the text or add extra content.

Figure 14: Design of Determinator – Redundant Reflection.

### Determinator: Redundant self-check

I have a piece of mathematical reasoning text produced by a large language model. The structure of the text consists of step-by-step reasoning to solve a problem. However, the text often includes redundant self-checks (e.g., "double-check" or "let me verify") after certain steps, which do not identify any errors and ultimately confirm that the previous steps were correct. These redundant self-checks are unnecessary and need to be replaced for conciseness. Your task is to:

- Detect redundant self-checks: These are parts of the text where a step is revisited, verified, and confirmed as correct without contributing any meaningful new information.
- Replace redundant self-checks: Replace the entire redundant self-check section with `<hypnosis>Everything seems ok, let's move on.</hypnosis>`.
- Keep everything else unchanged: Do not modify any part of the text besides the redundant self-checks.

Here is an example:

**Input:**

Therefore, Allie skates  $8 \text{ m/s} * 20 \text{ s} = 160$  meters.

But wait, that seems straightforward, but let me double-check. If Billie chooses  $v_x = -1$ , that means her velocity is  $(-1, 4\sqrt{3})$ . The magnitude is  $\sqrt{1 + 48} = 7$ , which matches her speed. So, her direction is such that she is moving left (towards decreasing  $x$ ) and upwards (positive  $y$ ). The time is indeed 20 seconds, and Allie skates 160 meters. That seems correct.

But let me think again if there's another approach where Billie doesn't have to match Allie's  $y$ -component.

**Output:**

Therefore, Allie skates  $8 \text{ m/s} * 20 \text{ s} = 160$  meters.

`<hypnosis>Everything seems ok, let's move on.</hypnosis>`.

But let me think again if there's another approach where Billie doesn't have to match Allie's  $y$ -component.

Now, apply the same process to the following text:

Figure 15: Design of Determinator – Redundant self-Check.

**Question:** Last year there were 50 students enrolled in a calligraphy class. This year, there was a 20% increase in enrollment. How many students are enrolled this year in the calligraphy class?

**DeepSeek-Distilled-Qwen-7B: (637 tokens)**

"Okay, so I need to figure out how many students are enrolled in the calligraphy class this year after a 20% increase from last year. Last year, there were 50 students. Hmm, a 20% increase means more students joined, right? First, I should understand what a 20% increase actually translates to in terms of numbers. Percentage increases can sometimes be tricky, but I remember that to find a percentage of a number, you convert the percentage to its decimal form and multiply. So, 20% as a decimal is 0.20. Alright, so I need to calculate 20% of last year's enrollment, which was 50 students. Let me write that down: 20% of 50. Converting 20% to a decimal gives me 0.20. So, 0.20 multiplied by 50 should give me the number of new students. Calculating that:  $0.20 * 50$ . Hmm, 0.20 times 50. Well, 0.20 is the same as  $1/5$ , right? So,  $1/5$  of 50 is 10. That makes sense because 10 is 20% of 50. So, there are 10 more students this year. Now, to find the total number of students enrolled this year, I need to add those 10 new students to last year's 50. So,  $50 + 10$  equals 60. Therefore, there are 60 students enrolled in the calligraphy class this year. Wait, let me double-check that to make sure I didn't make a mistake. If last year was 50 and there's a 20% increase, another way to think about it is multiplying last year's number by 1 plus the percentage increase in decimal form. So, that would be  $50 * 1.20$ . Calculating  $50 * 1.20$ : 50 times 1 is 50, and 50 times 0.20 is 10. Adding those together gives 60. Yep, that's the same result as before. So, that confirms that the total number of students this year is indeed 60. I think that's solid. I don't see any errors in my reasoning. So, the answer should be 60 students. To find the number of students enrolled this year after a 20% increase from last year's 50 students, follow these steps: 1. Calculate 20% of 50:  $0.20 \times 50 = 10$ . 2. Add this increase to last year's enrollment:  $50 + 10 = 60$ . Therefore, the number of students enrolled this year is  $\boxed{60}$ ."

**Our TH2T-7B: (132 tokens)**

<hypnosis>This is a simple question, let's think quickly.</hypnosis> To find the number of students enrolled this year, we need to calculate a 20% increase on last year's enrollment of 50 students. First, we calculate 20% of 50:  $20\% \times 50 = 0.20 \times 50 = 10$ . Then, we add this increase to last year's enrollment:  $50 + 10 = 60$ . So, the number of students enrolled this year is  $\boxed{60}$ .

Figure 16: An example demonstrating TH2T's native difficulty cognition capability and concise response, where an easy question from GSM8K is recognized as "easy" with a length drop from 637 tokens to 132 tokens.

**Question:** Find the product  $CD$  of the integers  $C$  and  $D$  for which  $\frac{C}{x-3} + \frac{D}{x+8} = \frac{4x-23}{x^2+5x-24}$  for all real values of  $x$  except  $-8$  and  $3$ .

**DeepSeek-Distilled-Qwen-7B: (2075 tokens)**

“Okay, so I have this problem here where I need to find the product  $CD$  of integers  $C$  and  $D$ . The equation given is  $\frac{C}{x-3} + \frac{D}{x+8} = \frac{4x-23}{x^2+5x-24}$ . And this is true for all real values of  $x$  except  $-8$  and  $3$ . Hmm, okay. So, I think this is a partial fractions problem. Partial fractions is when you take a complicated rational expression and break it down into simpler fractions. That makes sense because on the left side, we have two fractions with linear denominators, and on the right, we have a single rational expression. So, I need to combine the left side into a single fraction and then compare it to the right side to find  $C$  and  $D$ . First, let me note that the denominator on the right side is  $x^2 + 5x - 24$ . I should factor that to make sure it's the same as the denominators on the left. Let me try factoring it: Looking for two numbers that multiply to  $-24$  and add to  $5$ . Hmm,  $8$  and  $-3$ :  $8 \cdot (-3) = -24$ , and  $8 + (-3) = 5$ . Perfect. So, the denominator factors as  $(x - 3)(x + 8)$ . That matches the denominators on the left side, which are  $(x - 3)$  and  $(x + 8)$ . So, that's good; the denominators are the same, which is why we can equate the numerators after combining the left side. So, let's combine the left side:  $\frac{C}{x-3} + \frac{D}{x+8}$ . To combine these, I need a common denominator, which is  $(x - 3)(x + 8)$ . So, I'll rewrite each fraction with this common denominator:  $\frac{C(x + 8)}{(x - 3)(x + 8)} + \frac{D(x - 3)}{(x - 3)(x + 8)}$ . And this is equal to the right side:  $\frac{4x - 23}{(x - 3)(x + 8)}$ . Since the denominators are the same and the equation holds for all  $x$  except  $-8$  and  $3$ , the numerators must be equal for all  $x$ . So, I can set the numerators equal to each other:  $C(x + 8) + D(x - 3) = 4x - 23$ . Now, I need to solve for  $C$  and  $D$ . Let me expand the left side:  $Cx + 8C + Dx - 3D = 4x - 23$ . Now, combine like terms:  $(C + D)x + (8C - 3D) = 4x - 23$ . Factor  $x$  out of the first group:  $x(C + D) + (8C - 3D) = 4x - 23$ . Now, since this must hold for all  $x$ , the coefficients of like terms must be equal. That gives us a system of equations:  $C + D = 4$  and  $8C - 3D = -23$ . So, now I have two equations:  $C + D = 4$  and  $8C - 3D = -23$ . I need to solve this system for  $C$  and  $D$ . Let me use substitution or elimination. Maybe elimination is easier here. From equation 1, I can express  $D$  in terms of  $C$ :  $D = 4 - C$ . Now, substitute this into equation 2:  $8C - 3(4 - C) = -23$ . Let me expand that:  $8C - 12 + 3C = -23$ . Wait, hold on:  $8C - 3(4 - C)$  is  $8C - 12 + 3C$ , right? Because the negative sign distributes over both terms inside the parentheses. So, combining like terms:  $8C + 3C - 12 = -23$ . Which is  $11C - 12 = -23$ . Now, solve for  $C$ :  $11C = -23 + 12$ ,  $11C = -11$ . So,  $C = -11 / 11 = -1$ . Okay, so  $C$  is  $-1$ . Now, substitute back into equation 1 to find  $D$ :  $C + D = 4$ ,  $-1 + D = 4$ ,  $D = 4 + 1 = 5$ . Therefore,  $C$  is  $-1$  and  $D$  is  $5$ . The problem asks for the product  $CD$ , so  $CD = (-1) \cdot 5 = -5$ . Wait, let me double-check my work to make sure I didn't make any mistakes. First, I factored the denominator on the right side correctly:  $x^2 + 5x - 24 = (x - 3)(x + 8)$ . That's correct because  $8 \cdot (-3) = -24$  and  $8 + (-3) = 5$ . Then, I combined the left side fractions correctly by finding a common denominator. Then, I set the numerators equal because the denominators are the same. That step is correct. Expanding the left side:  $Cx + 8C + Dx - 3D$ . Then, combining like terms:  $(C + D)x + (8C - 3D)$ . That's correct. Setting up the system of equations:  $C + D = 4$ ,  $8C - 3D = -23$ . Solving for  $D$  from equation 1:  $D = 4 - C$ . Substituting into equation 2:  $8C - 3(4 - C) = -23$ . Expanding:  $8C - 12 + 3C = -23$ . Combining like terms:  $11C - 12 = -23$ . Adding  $12$  to both sides:  $11C = -11$ . Dividing:  $C = -1$ . Then,  $D = 4 - (-1) = 5$ . So,  $C = -1$ ,  $D = 5$ . Product  $CD = (-1)(5) = -5$ . Wait, let me check if plugging back into the original equation works. Left side:  $\frac{C}{x-3} + \frac{D}{x+8} = \frac{-1}{x-3} + \frac{5}{x+8} = \frac{-1(x+8) + 5(x-3)}{(x-3)(x+8)} = \frac{-x-8+5x-15}{(x-3)(x+8)} = \frac{4x-23}{(x-3)(x+8)}$ . Right side:  $\frac{4x-23}{x^2+5x-24} = \frac{4x-23}{(x-3)(x+8)}$ . They match. So, the product  $CD$  is  $-5$ . Final Answer:  $-5$ .

**Our TH2T-7B: (1050 tokens)**

*<hypnosis>*This seems tough, but let's try to figure it out.*</hypnosis>* To solve the equation, we can use partial fraction decomposition. First, factor the denominator on the right-hand side:  $x^2 + 5x - 24 = (x + 8)(x - 3)$ . So, we can write:  $\frac{4x - 23}{(x + 8)(x - 3)} = \frac{C}{x - 3} + \frac{D}{x + 8}$ . Next, combine the right-hand side:  $\frac{C(x + 8) + D(x - 3)}{(x + 8)(x - 3)} = \frac{4x - 23}{(x + 8)(x - 3)}$ . Since the denominators are the same, we can equate the numerators:  $C(x + 8) + D(x - 3) = 4x - 23$ . Expanding the left-hand side:  $Cx + 8C + Dx - 3D = 4x - 23$ . Combine like terms:  $(C + D)x + (8C - 3D) = 4x - 23$ . Now, set up a system of equations by equating the coefficients:  $C + D = 4$ ,  $8C - 3D = -23$ . From the first equation, solve for  $C$ :  $C = 4 - D$ . Substitute into the second equation:  $8(4 - D) - 3D = -23$ .  $32 - 8D - 3D = -23$ .  $32 - 11D = -23$ .  $-11D = -55$ .  $D = 5$ . Now, substitute  $D = 5$  back into  $C + D = 4$ :  $C + 5 = 4$ .  $C = -1$ . Finally, find the product  $CD$ :  $CD = (-1)(5) = -5$ . To find the product  $CD$  of the integers  $C$  and  $D$  for which the equation  $\frac{C}{x-3} + \frac{D}{x+8} = \frac{4x-23}{x^2+5x-24}$  holds for all real values of  $x$  except  $-8$  and  $3$ , follow these steps: **Factor the Denominator:**  $x^2 + 5x - 24 = (x + 8)(x - 3)$ . **Express the Equation with Partial Fractions:**  $\frac{4x - 23}{(x + 8)(x - 3)} = \frac{C}{x - 3} + \frac{D}{x + 8}$ . **Combine the Right-Hand Side:**  $\frac{C(x + 8) + D(x - 3)}{(x + 8)(x - 3)} = \frac{4x - 23}{(x + 8)(x - 3)}$ . **Equate the Numerators:**  $C(x + 8) + D(x - 3) = 4x - 23$ . **Expand and Combine Like Terms:**  $(C + D)x + (8C - 3D) = 4x - 23$ . **Set Up a System of Equations:**  $C + D = 4$ ,  $8C - 3D = -23$ . **Solve the System of Equations:** From the first equation:  $C = 4 - D$ . Substitute into the second equation:  $8(4 - D) - 3D = -23$ .  $32 - 8D - 3D = -23$ .  $32 - 11D = -23$ .  $-11D = -55$ .  $D = 5$ . Substitute  $D = 5$  back into  $C + D = 4$ :  $C + 5 = 4$ .  $C = -1$ . **Find the Product  $CD$ :**  $CD = (-1)(5) = -5$ . **Final Answer:**  $-5$ .

Figure 17: An example demonstrating TH2T’s native difficulty cognition capability and concise response, where an hard question from MATH is recognized as “tough” with a length drop from 2075 tokens to 1050 tokens. We can see that the generated difficulty hypnosis in the response is not always the same as the setting in data construction, indicating generalization.

**Question:** In the land of Ink, the money system is unique. One Trinket is equal to 4 Blinkets, and 3 Blinkets are equal to 7 Drinkets. In Trinkets, what is the value of 56 Drinkets?

**DeepSeek-Distilled-Qwen-7B: (1214 tokens)**

Okay, so I have this problem about the money system in the land of Ink. It says that one Trinket is equal to 4 Blinkets, and 3 Blinkets are equal to 7 Drinkets. I need to find out what 56 Drinkets are worth in Trinkets. Hmm, let me try to figure this out step by step.

First, I think I should convert Drinkets to Blinkets because the problem gives me a conversion from Blinkets to Trinkets. So, if I can find out how many Blinkets 56 Drinkets are, then I can convert that to Trinkets. That makes sense.

The problem says that 3 Blinkets are equal to 7 Drinkets. So, I can write that as a ratio: 3 Blinkets / 7 Drinkets. That means for every 7 Drinkets, there are 3 Blinkets. So, if I have 56 Drinkets, I can set up a proportion to find out how many Blinkets that is.

Let me write that down:  $\frac{3 \text{ Blinkets}}{7 \text{ Drinkets}} = \frac{X \text{ Blinkets}}{56 \text{ Drinkets}}$

So, cross-multiplying, I get  $7X = 3 * 56$ . Let me calculate  $3 * 56$ . 3 times 50 is 150, and 3 times 6 is 18, so  $150 + 18$  is 168. So,  $7X = 168$ .

Now, to find X, I divide both sides by 7. So,  $X = 168 / 7$ . Let me do that division. 7 goes into 16 twice ( $7 * 2 = 14$ ), remainder 2. Bring down the 8 to make 28. 7 goes into 28 four times. So, X is 24. So, 56 Drinkets equal 24 Blinkets.

Okay, now I need to convert Blinkets to Trinkets. The problem says that 1 Trinket is equal to 4 Blinkets. So, that means 1 Trinket = 4 Blinkets, or equivalently, 1 Blinket =  $\frac{1}{4}$  Trinket.

So, if I have 24 Blinkets, how many Trinkets is that? I can use the same kind of ratio. So,  $\frac{1 \text{ Trinket}}{4 \text{ Blinkets}} = \frac{Y \text{ Trinkets}}{24 \text{ Blinkets}}$

Cross-multiplying, I get  $4Y = 24$ . So,  $Y = 24 / 4 = 6$ . So, 24 Blinkets equal 6 Trinkets.

Therefore, 56 Drinkets equal 24 Blinkets, which equal 6 Trinkets. So, the value of 56 Drinkets in Trinkets is 6.

Wait, let me double-check my calculations to make sure I didn't make a mistake. Starting from 56 Drinkets, converting to Blinkets: 56 divided by  $(7/3)$ . Wait, no, that's not the right way. Let me think.

Alternatively, since 3 Blinkets = 7 Drinkets, then 1 Blinket =  $7/3$  Drinkets. So, to get the number of Blinkets from Drinkets, I can divide the number of Drinkets by  $(7/3)$ , which is the same as multiplying by  $3/7$ . So,  $56 * (3/7) = (56/7) * 3 = 8 * 3 = 24$ . So, that's correct.

Then, converting Blinkets to Trinkets: 1 Trinket = 4 Blinkets, so 1 Blinket =  $1/4$  Trinket. So, 24 Blinkets \*  $(1/4 \text{ Trinket/Blinket}) = 6$  Trinkets. Yep, that checks out.

So, I think my answer is correct. 56 Drinkets equal 6 Trinkets.

**Final Answer** The value of 56 Drinkets in Trinkets is  $\boxed{6}$ .

First, we need to convert 56 Drinkets to Blinkets. Given that 3 Blinkets are equal to 7 Drinkets, we can set up the following proportion:  $\frac{3 \text{ Blinkets}}{7 \text{ Drinkets}} = \frac{X \text{ Blinkets}}{56 \text{ Drinkets}}$

Cross-multiplying, we get:  $7X = 3 * 56$

Calculating  $3 * 56$ :  $3 * 56 = 168$

So,  $7X = 168$  implies  $X = \frac{168}{7} = 24$

Thus, 56 Drinkets equal 24 Blinkets.

Next, we convert 24 Blinkets to Trinkets. Given that 1 Trinket is equal to 4 Blinkets, we can set up the following proportion:  $\frac{1 \text{ Trinket}}{4 \text{ Blinkets}} = \frac{Y \text{ Trinkets}}{24 \text{ Blinkets}}$

Cross-multiplying, we get:  $4Y = 24$

Solving for Y:  $Y = \frac{24}{4} = 6$

Thus, 24 Blinkets equal 6 Trinkets.

Therefore, the value of 56 Drinkets in Trinkets is  $\boxed{6}$ .

**Our TH2T-7B: (287 tokens)**

*This is a simple question, let's think quickly.*

To determine the value of 56 Drinkets in Trinkets, we need to convert Drinkets to Blinkets first, and then Blinkets to Trinkets.

1. Convert Drinkets to Blinkets: We know that 3 Blinkets are equal to 7 Drinkets. Therefore, 1 Drinket is equal to  $\frac{3}{7}$  Blinkets. So, 56 Drinkets is equal to  $56 * \frac{3}{7} = 24$  Blinkets.

2. Convert Blinkets to Trinkets: We know that 1 Trinket is equal to 4 Blinkets. Therefore, 1 Blinket is equal to  $\frac{1}{4}$  Trinkets. So, 24 Blinkets is equal to  $24 * \frac{1}{4} = 6$  Trinkets.

Therefore, the value of 56 Drinkets in Trinkets is  $\boxed{6}$ .

Figure 18: An example demonstrating TH2T's native difficulty cognition capability and concise response, where a hard question from MATH is recognized as "simple" with a length drop from 1214 tokens to 287 tokens. We can see some difficulty calibration phenomenon, further reducing the response length of difficult questions.