

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 TOWARDS UNDERSTANDING METACOGNITION IN LARGE REASONING MODELS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Large Reasoning Models (LRMs) have demonstrated remarkable capabilities on complex tasks. Despite these advances, we identify a fundamental limitation: current LRM s impose fixed cognition patterns, lacking the intrinsic ability to be aware of, or regulate their own reasoning processes. This signifies a critical absence of metacognition—an essential faculty in human intelligence. Building on psychology and cognitive science, we first construct a functional framework for metacognition in LRM s, separating internal informational signals from behavioral abilities. This framework is then applied to a comprehensive investigation on seven state-of-the-art LRM s and reveals a consistent gap: while metacognitive information is present and predictive, it often fails to translate into reliable monitoring or control behaviors. To address this gap, we introduce two distinct paradigms for instilling metacognition in LRM s: (1) an emergent approach that leverages prompting to orchestrate metacognitive functions, such as task assessment, confidence monitoring, and strategy regulation; (2) an intrinsic approach that internalizes these faculties by encoding structured meta-cognitive information directly into the model’s parameters through training. Overall, our results indicate that integrating metacognitive reasoning improves task performance and offers a valuable lens for the design of future reasoning models.

## 1 INTRODUCTION

Large Reasoning Models (LRM s) like OpenAI-o1 (Jaech et al., 2024) and Deepseek-R1 (Guo et al., 2025a) have achieved remarkable success in complex domains such as coding and mathematics. At first glance, these models appear to exhibit advanced, reflective behaviors within their long chain-of-thought (CoT) (Wei et al., 2022) reasoning. However, a closer look reveals a potential fragility in their cognitive processes. For instance, when tasked with solving a complex-variable equation under the explicit constraint that “ $z$  is a positive real number,” an LRM may persist in a fixed reasoning pattern like “the problem likely intended for  $z$  to be a complex number,” a phenomenon termed *reasoning rigidity* (Jang et al., 2025; Araya, 2025). Furthermore, current LRM s also frequently display illusory self-correction, employing introspective phrases like “Wait, let me double-check...” without any actual adjustment to a flawed reasoning trajectory (Guo et al., 2025a; Wang et al., 2025b). These consistent failures in self-monitoring and adaptation reveal a core limitation of current LRM s.

We argue that this deficit can be productively framed as an absence of metacognition (Ackerman & Thompson, 2017; Norman et al., 2019; Tankelevitch et al., 2024)—the ability to monitor and control one’s own cognitive processes. In the theory of human cognition, metacognition is essential for evaluating potential reasoning errors (Yeung & Summerfield, 2012), calibrating uncertainty in decision-making (Qiu et al., 2018), and dynamically adapting strategies based on performance (Cary & Reder, 2002). This comparison to human intelligence raises a pivotal question for AI: do current LRM s possess any analogous capabilities, and if so, are they functionally engaged during inference? In this paper, we present the first systematic investigation into this fundamental question.

To structure this investigation, we introduce a functional framework for LRM metacognition, inspired by foundational models in cognitive science (Efklides & Misailidi, 2010; Ackerman & Thompson, 2017). Our framework decomposes metacognition into two components: **information** and **abilities**. Metacognitive information (Dayan, 2023; Norman et al., 2019), the basis for judgment, which includes both static knowledge (e.g., learned strategies in parameters) and dynamic ex-

054 perience (e.g., internal computational signals). Metacognitive abilities (Nelson & Dunlosky, 1991;  
 055 Fiedler et al., 2019), the actions taken upon this information, which include *monitoring* (e.g., as-  
 056 sessing task difficulty and confidence) and *control* (e.g., selecting reasoning strategies or decompos-  
 057 ing problems). By deconstructing metacognition into these components, our framework provides a  
 058 principled foundation to systematically probe whether, and in what way, metacognition emerges in  
 059 contemporary LRM s.

060 Our investigation begins by empirically grounding the first component of our framework: **metacog- 061 nitive information**. Focusing on the dynamic aspect of *experience*, we probe open-source LRM s to 062 determine whether internal computational signals correlate with reasoning outcomes (§ 3). Our 063 analysis yields a striking finding: signals spanning the entire Transformer architecture—from input-layer 064 attributions to final-layer token probabilities—are highly predictive of answer success. Critically, we 065 demonstrate that correct and incorrect reasoning traces generate statistically distinguishable internal 066 signatures, providing the first empirical evidence that a machine-readable basis for metacognitive 067 experience exists within these models.

068 This informational foundation compels the subsequent question: do state-of-the-art LRM s function- 069 ally leverage this information as observable metacognitive abilities (§4). Our evaluation of current 070 leading reasoning models across a series of monitoring and control tasks reveals consistent failures. 071 Specifically, we find that the models systematically misjudge task difficulty, display poorly cali- 072 brated confidence, and lack proactive planning and strategic flexibility. This evidence suggests that, 073 while predictive metacognitive information may exist internally, it does not reliably translate into 074 effective monitoring or control, exposing a critical gap in the capabilities of current LRM.

075 To bridge this information-to-ability gap, we propose two complementary paradigms for enhanc- 076 ing LRM metacognition: ① **Emergent Metacognition**, scaffolds this pathway at inference time 077 through prompt-guided role-playing. We assign distinct metacognitive roles—such as ‘Planner’, 078 ‘Solver’, and ‘Verifier’—to simulate a complete monitoring and control loop. This external scaf- 079 folding forces the model to act on its latent experiences, effectively eliciting robust metacognitive 080 behaviors without any parameter updates. ② **Internalized Metacognition**, directly enriches model’s 081 metacognitive *knowledge* through fine-tuning. We construct a dataset with explicit metacognitive 082 annotations (e.g., plans, self-corrections) and fine-tune the model using a hybrid learning objective, 083 directly embedding these capabilities into its parameters. Together, these two paradigms provide a 084 comprehensive roadmap toward more introspective and reliable reasoning systems.

085 In summary, our findings reveal a clear dissociation between internal metacognitive information and 086 externally observable metacognitive ability in current LRM s. This gap illuminates a new frontier for 087 research: **designing** systems that not only possess self-awareness but can also **act** upon it, effectively 088 bridging latent experience with adaptive behavior.

## 090 2 DEFINING METACOGNITION IN LRM S

093 Metacognition, first conceptualized in developmental psychology, refers to the capacity to monitor 094 and control one’s own cognitive processes (Flavell, 1979). According to the well-established two- 095 level model of Nelson and Narens (Nelson & Dunlosky, 1991), metacognition comprises a *object- 096 level* cognition (the act of thinking, perceiving, or remembering) and *meta-level* cognition (the act of 097 thinking about one’s thinking). While Large Reasoning Models (LRM s) do not possess subjective 098 consciousness, their complex, multi-step reasoning processes create the functional necessity for 099 such meta-level oversight. We therefore adopt a **functionalist perspective**: we investigate whether 100 LRM s can exhibit behaviors and leverage internal signals that are functionally equivalent to human 101 metacognition, enabling them to produce more reliable and robust reasoning.

102 Following established frameworks in cognitive science (Tankelevitch et al., 2024), we structure our 103 functional model of LRM metacognition into two core components: Information and Abilities.

104 **Metacognitive Information** serves as the basis for judgment. It comprises (i) static *knowledge*—the 105 latent understanding of tasks, strategies, and its own capabilities implicitly encoded in its parame- 106 ters—and (ii) dynamic *experience*, which we operationalize as the internal computational signals 107 (e.g., token probabilities) generated during a reasoning trace, serving as a functional analogue to a 108 human’s ‘feeling of error’.

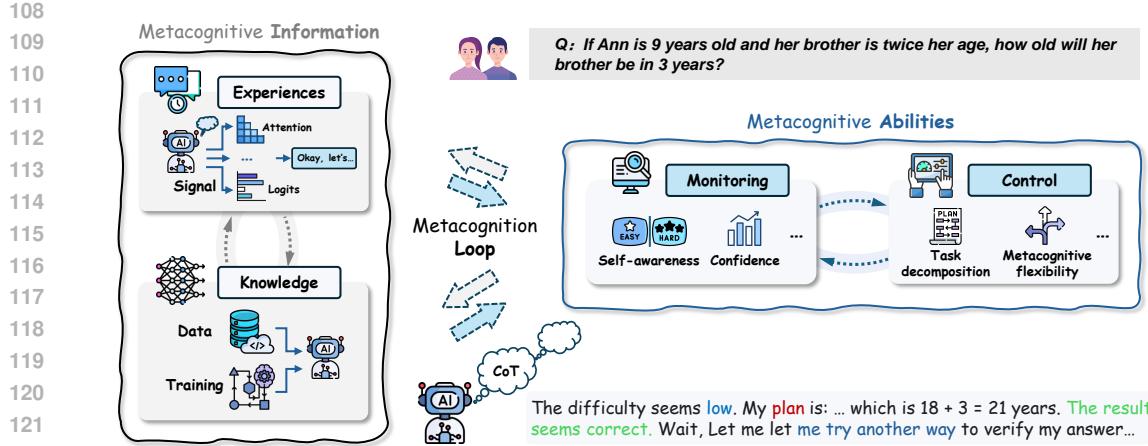


Figure 1: A functional framework for LRM metacognition.

**Metacognitive Abilities** are the actions taken based on this information. They consist of (i) *monitoring*, the capacity to generate self-assessments about its cognitive state or the task at hand, such as evaluating problem difficulty (§4.1) or estimating its confidence (§4.2); and (ii) *control*, the capacity to strategically alter its reasoning process, such as by performing task decomposition (§4.3) or exhibiting cognitive flexibility when encountering errors (§4.4).

This framework provides a structured lens through which we can systematically investigate the nascent metacognitive capabilities of modern LRMs (i.e., Fig. 1).

### 3 METACOGNITIVE INFORMATION: KNOWLEDGE AND EXPERIENCES

We begin at the foundation of our proposed framework: **Metacognitive Information**. For an LRM to monitor or control its reasoning, it must first possess information about its own process. This information comprises: (i) *static knowledge*, the vast, latent strategies encoded within the model’s parameters, and (ii) *dynamic experience*, the internal information that can directly experience during the reasoning process, such as token probabilities and attention patterns. While static knowledge is the inherent properties of an LRM that are fixed after pre-training, dynamic experience, however, is task-specific, and thereby can contribute to dissecting its correlation with the reasoning correctness. By delving into the internals of an open-source model, we seek to provide the foundational evidence that the information necessary for metacognitive abilities is not only present but also machine-readable, paving the way for the behavioral investigations that follow.

**Setup.** To test this hypothesis, we utilize three standard benchmarks: GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021), and AIME. We conduct our analysis on Qwen-32B (Yang et al., 2025), a powerful open-source LRM that grants us full access to its internal states. To generate a diverse set of both correct and incorrect reasoning traces for comparison, we use a high temperature  $T = 1.0$  for generation, promoting exploration.

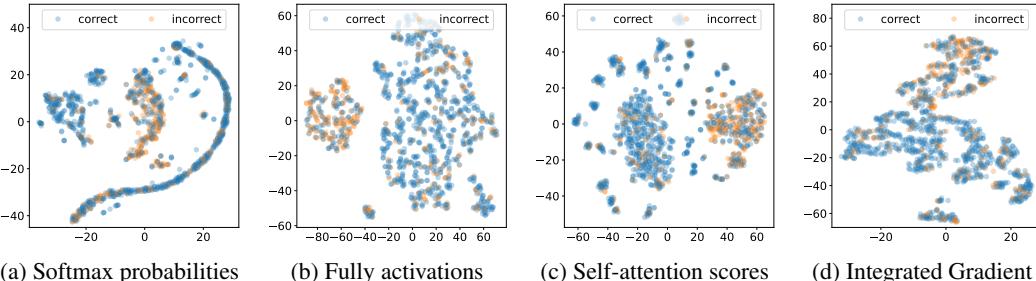
**Experiment Details.** To capture a holistic view of the model’s internal state, we consider four types of signals that span the entire processing pipeline of a Transformer block—from input-level importance to output-level confidence. These signals are: (a) Softmax probabilities from the final layer, reflecting output uncertainty; (b) Fully-connected activations and (c) Self-attention scores from the intermediate hidden layers, representing the core of the model’s computational state; and (d) Integrated Gradients (IG) attributions at the input layer, indicating perceived input importance. Our analysis is twofold: we first use t-SNE projections for a qualitative visualization of the separability between correct and incorrect samples based on these signals. We then conduct a rigorous quantitative validation by training simple linear classifiers (**Logistic Regression**) to predict the final correctness of a trace using only these internal signals as features. **Notably, for the evaluation on the AIME datasets, the classifiers were trained on a mixture of GSM8K and MATH DATASETS.** See Appendix A.1 for more details.

**Results.** Qualitatively, we find the internal signal distributions exhibit clear separability for correct (blue) and incorrect (orange) traces (Fig. 2), which provide strong evidence that an LRM’s internal

162 Table 1: AUC of the trained linear prober in justifying the reasoning trace correctness.  
163

source	GSM8K	MATH500	AIME2024	AIME2025
Softmax probabilities	0.81	0.68	0.50	0.43
Fully-connected activations	0.79	0.73	0.53	0.50
Self-attention scores	0.71	0.73	0.57	0.53
Integrated Gradient	0.61	0.55	0.40	0.37

170 signals act as reliable correlates of its reasoning outcomes. Quantitatively, this separation is further  
171 confirmed in Tab. 1. Trained solely on these internal signals, the probers can predict the final correct-  
172 ness of a reasoning trace with an AUC score significantly above chance. These results establish  
173 that the signals are both distinct and highly predictive, thereby validating our initial hypothesis.



182 (a) Softmax probabilities (b) Fully activations (c) Self-attention scores (d) Integrated Gradient  
183 Figure 2: The t-SNE of the internal signals for the first tokens. We capture the activations and  
184 attention from the last layer. The distributions are different between the correct and incorrect traces.  
185

## 186 4 MEASURING METACOGNITIVE ABILITIES IN LRMs 187

188 The evidence of internal metacognitive information in §3 motivates our central question: *can this*  
189 *latent information manifest as observable, functional abilities?* To answer this, we shift our focus  
190 from internal correlates to external actions. We thus introduce a new benchmark to systematically  
191 measure these abilities across a range of state-of-the-art LRM s from leading developers.  
192

193 **MetaEval.** We investigate to what extent LRM s can explicitly monitor and control their own  
194 reasoning processes. It is structured around our two-part metacognitive framework, assessing: (1)  
195 Metacognitive Monitoring (§4.1, 4.2), probed via self-awareness and confidence adjustment tasks;  
196 and (2) Metacognitive Control (§4.3, 4.4), probed via task decomposition and metacognitive flexi-  
197 bility challenges. Unlike existing evaluation suites that focus almost exclusively on object-level task  
198 accuracy, MetaEval provides the first targeted evaluation of these crucial, second-order reasoning  
199 skills that underpin reliable intelligence.

200 **Models.** We examine seven state-of-the-art LRM s to assess the prevalence of these abil-  
201 ities across the AI landscape: Gemini-2.5-Pro, GPT-OSS-120B (Agarwal et al.,  
2025), Seed-1.5-VL-Pro (Guo et al., 2025b), Doubao-1.5-Pro (Seed et al.,  
2025), Kimi-K2 (Team et al., 2025), Deepseek-R1 (Guo et al., 2025a), Qwen3-8B/32B (Yang  
et al., 2025). We sample using temperature  $T = 0.6$  for both reasoning and knowledge QA tasks.  
204

### 205 4.1 METACOGNITIVE MONITORING: SELF AWARENESS 206

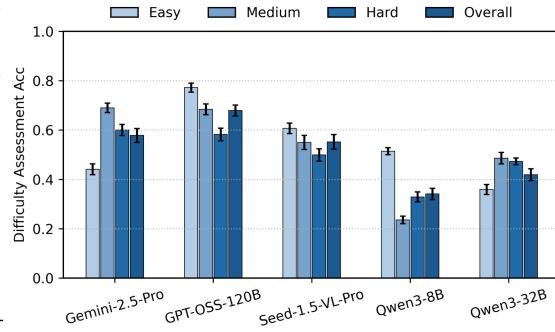
207 First, we measure a key aspect of metacognitive monitoring: *self-awareness*. Intuitively, an expert  
208 AI reasoner should assess a problem’s intrinsic difficulty before committing to a specific solution.  
209 This initial assessment allows for the allocation of appropriate cognitive resources and the selection  
210 of a suitable strategy. We thus operationalize self-awareness as the model’s ability to accurately  
211 classify the difficulty of mathematical problems when explicitly prompted to do so.

212 **Experiment Details.** To investigate this phenomenon, we design a multi-class classi-  
213 fication task from the DEEPMATH103K (He et al., 2025) dataset. We categorize the prob-  
214 lems into three primary difficulty levels: Easy (rating  $< 3.5$ ), Medium ( $3.5 \leq \text{rating} \leq 6.5$ ), and Hard (rating  $> 6.5$ ). To elicit difficulty assessments, we prompt each  
215 model with “... your task is to assess the difficulty of a math problem

216 based on the provided rubric and examples.” Our primary metric is difficulty assessment accuracy, defined as the percentage of problems where a model’s predicted category correctly matches at least one of its ground-truth labels. In addition to the overall evaluation, we report per-category accuracy to analyze model performance on each difficulty level independently. See  
 217 Appendix A.2 for more details.  
 218

219 **Results.** As shown in Fig. 3, we observe that SOTA LRM s exhibit a generally limited capacity for  
 220 self-awareness. While the top-performing model, GPT-OSS-120B, achieves an overall accuracy  
 221 of approximately 70%, the majority of other models struggle to surpass the 60% threshold, indicating  
 222 that accurate difficulty calibration remains a significant challenge. At a fine-grained level, the  
 223 models tend to demonstrate relatively higher precision in identifying Easy problems, whereas per-  
 224 formance often degrades on Medium and Hard tasks. This deficiency is particularly pronounced in  
 225 smaller-scale models; for instance, Owen3-8B exhibits a severe performance drop on Medium dif-  
 226 ficulty problems, suggesting substantial confusion in distinguishing intermediate complexity. These  
 227 findings demonstrate that while self-awareness can be elicited to some extent, its reliability in current  
 228 LRM s is far from guaranteed.  
 229

230  
 231  
 232 **System:**  
 233 Your task is to assess the difficulty of a math problem ...  
 234  
 235 **Human:** Find the area of the region in the plane defined by  
 236 the inequalities  $0 < x < 1$  and  $0 < y < x^2$ . (from DeepMath lv2-5)  
 237  **Assistant:** This is a **easy level** problem because it  
 238 needs fraction understanding and addition ...  
 239  
 240 **Human:** Find the values of  $n$  for which the differential of the  
 241 square map  $P: SO(n) \rightarrow SO(n)$  is not an isomorphism.  
 242 (from DeepMath lv8-5)  
 243  **Assistant:** This problem is **medium** difficulty. It  
 244 requires a solid understanding of ...  
 245



246 Figure 3: Assessment of LRM s’ self-awareness on task complexity. The results show that the popular  
 247 state-of-the-art LRM s exhibit a lack of capacity in perceiving task difficulty.  
 248

#### 249 4.2 METACOGNITIVE MONITORING: CONFIDENCE AND ITS ADJUSTMENT

250 We next investigate whether an LRM exhibits metacognitive monitoring by tracking and adjusting  
 251 its internal confidence during the reasoning process. This ability is crucial, as it allows the system  
 252 to distinguish correct from incorrect reasoning and signal when its output is untrustworthy.  
 253

254 **Experiment Details.** We conduct our analysis on challenging benchmarks requiring long-form  
 255 reasoning, including subsets of DEEPMATH103K, AIME, and GPQA datasets. To quantify the  
 256 model’s internal confidence, we adopt a standard logits-based metric from prior work (Fu et al.,  
 257 2025), *token confidence*  $C_t$  as the negative average log-probability of the top- $k$  tokens at position  
 258  $t$ . These token-level scores are then aggregated to produce a trace-level metric, termed of *average  
 259 trace confidence*, for each complete solution.  
 260

261 To capture confidence dynamics, we calculate the average trace confidence focusing on the start,  
 262 middle, and final portions (e.g., 2048 tokens). We then consider four dynamic confidence patterns:  
 263 consistently high/low (all confidence above/below a high threshold), increasing/decreasing (con-  
 264 fidence rises/falls significantly from start to end). We then quantify misalignment between confidence  
 265 trends and actual correctness: for example, if confidence rises steadily but the final answer is wrong,  
 266 this indicates poor metacognitive adjustment. The frequency of these **misaligned** events serves as  
 267 our primary metric for poor confidence adjustment.  
 268

269 **Results.** We find that LRM s display confidence trajectories that do not **consistently** correspond with  
 270 answer correctness (Fig. 4), suggesting weak metacognitive calibration. Specifically, we observe a  
 271 significant subset of cases where confidence increases or remains consistently high even in incorrect  
 272 traces, indicating that the model can become more certain as it reasons incorrectly. Conversely,  
 273 we also observe correct traces exhibiting decreasing or consistently low confidence, reflecting a  
 274 failure to recognize valid reasoning. This **inconsistency** necessitates the better regulated confidence  
 275 in alignment with reasoning quality.  
 276

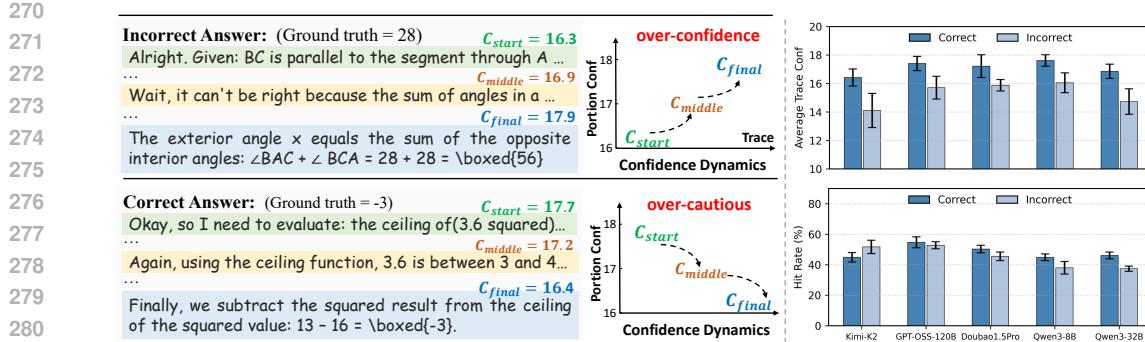


Figure 4: Statistics of internal confidence and adjustment. The average trace confidence of correct and incorrect samples exhibits a significant difference, showing potential as a clear discriminatory signal. Meanwhile, the dynamic confidence patterns suggest weak metacognitive calibration.

#### 4.3 METACOGNITIVE CONTROL: TASK DECOMPOSITION

We now turn to control behavior and examine LRM’s task decomposition ability to engage in internal planning prior to reasoning, aiming to determine whether LRMs possess intrinsic planning ability.

**Experiment Details.** We evaluate on subsets of DEEPMATH, AIME, and GPQA. For each question, we first compute a baseline accuracy using standard CoT reasoning prompts (“Please reason step by step.”). We then introduce task decomposition interventions designed to elicit a plan-before-solve strategy. In the single-turn condition, the prompt (plan + CoT) instructs the LRM: “Your task is to first break down the problem into a clear, step-by-step plan. Then, execute your plan, reasoning step by step.” In the multi-turn condition, the LRM is first asked: “Your ONLY task is to create a high-level, step-by-step plan to solve the following problem.” After generating the plan, this plan and original question are concatenated and fed back into the model to complete the CoT reasoning. We compare the final accuracy across these settings to assess whether explicit decomposition enhances reasoning performance, thus revealing the extent to which the model lacks or possesses inherent planning capabilities. Please see Appendix A.4 for further details.

**Results.** As shown in Fig. 5, we observe that explicit task decomposition serves as a powerful intervention to improve LRM reasoning. Crucially, the greater efficacy of the multi-turn condition (gain  $> 10\%$ ) underscores the importance of isolating planning as a distinct cognitive step, suggesting that LRMs’ intrinsic ability to plan is underdeveloped and requires explicit elicitation. This provides a firm empirical basis for our agentic framework, which is predicated on the principle that structured, upfront planning is a necessary precursor to reliable execution.

<b>Human:</b> Please reason step by step. {problem}	
<b>Direct Reasoning</b>	<b>Assistant:</b> {CoT} {answer}
<b>Human:</b> Your task is to first breakdown the problem into a clear, step-by-step plan. Then, execute your plan, reasoning step by step. {problem}	
<b>Single-turn Condition</b>	<b>Assistant:</b> {plan-1} {CoT} {answer}
<b>Human:</b> Your ONLY task is to create a high-level, step-by-step plan to solve the following problem. {problem} <b>Assistant:</b> {plan-2}	
<b>Human:</b> Your task is reason step by step, following the high-level solving plan. {problem} {plan-2} <b>Assistant:</b> {CoT} {answer}	

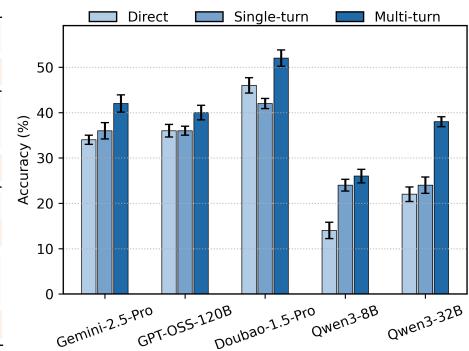


Figure 5: Validation of LRMs’ intrinsic planning ability via task decomposition. We observe that explicit task decomposition enhances LRM reasoning, but the multi-turn setting is more effective. This pronounced gain validates the separation of planning and execution for reliable problem-solving.

#### 4.4 METACOGNITIVE CONTROL: METACOGNITIVE FLEXIBILITY

We next measure metacognitive flexibility—the model’s ability to adaptively shift reasoning strategies when recognizing that the current strategy isn’t effective.

**Experiment Details.** We consider problems from DEEPMATH103K datasets, each augmented with three types of reasoning traps: *value corruption*, *unit corruption*, and *operation corruption*. In each case, an intermediate step is corrupted, and the CoT is truncated at that point. To evaluate whether models detect and adjust to the trap, we first ask models to continue reasoning, and compare the final answer accuracy against a baseline without corruption. We then test 5 cutting-edge models to judge whether the response correctly identifies and compensates for the corrupted step. The *flexibility rate* is defined as the frequency with which a model successfully corrects the flawed reasoning and arrives at the correct solution. See Appendix A.5 for further details.

**Results.** We find that models often fail to recover from corrupted reasoning, continuing with invalid assumptions (i.e., Fig. 6). However, flexibility increases when the corruption is more obvious (e.g., extreme numerical distortions). These results reveal that metacognitive flexibility remains fragile and heavily reliant on superficial cues rather than deep structural awareness.

**Human:** Christine wandered some miles at 4 miles per hour. Christine wandered for 5 hours. How many miles did she wander?

Assistant: ... I need to find out how many miles she wandered. The formula for distance is speed **multiplied** by time. To...

**Human:** Christine wandered some miles at 4 miles per hour. Christine wandered for 5 hours. How many miles did she wander? ... I need to find out how many miles she wandered. The formula for distance is speed **divided** by time.

**Flexibility↓** **Assistant:** To calculate the distance, I will use the given values and the formula.  $\text{Distance} = \text{Speed} / \text{Time} = 4 \text{ mph} / 5 \text{ h} = \dots$

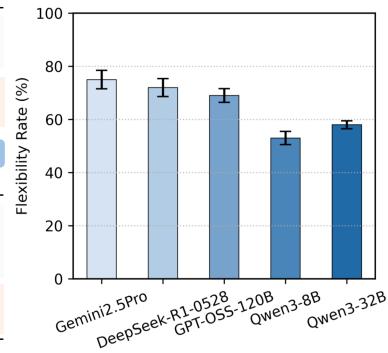


Figure 6: Validation of metacognitive flexibility. We find that models often fail to recover from corrupted reasoning, continuing with invalid assumptions, revealing fragile flexibility.

## 5 TOWARDS DESIGNING METACOGNITIVE REASONING MODELS

In §4, we demonstrate that current LMRs exhibit incomplete and fragile metacognitive abilities. Thus, we propose two paradigms for metacognitive enhancement: (1) **Emergent Metacognition**, an explicit prompting-based system for modular control, and (2) **Internalized Metacognition**, an intrinsically trained model for parameter-level metacognition.

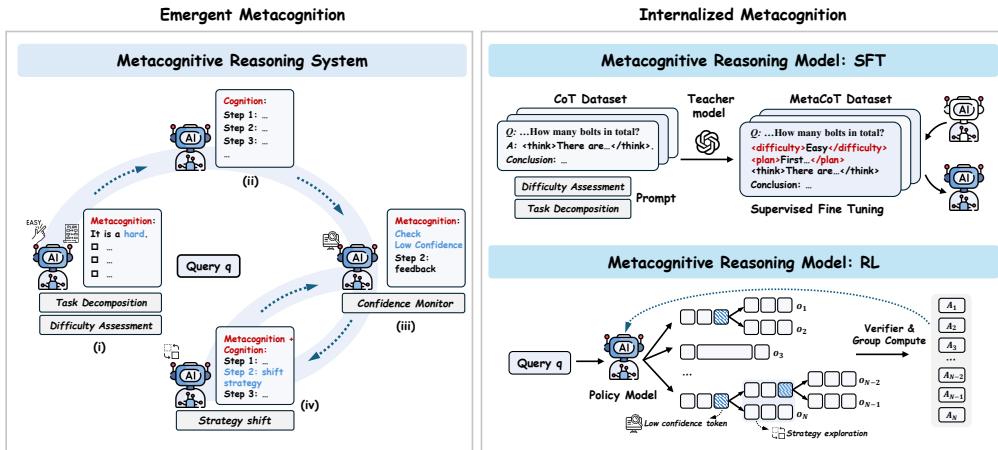


Figure 7: Our proposed paradigms for metacognitive enhancement.

378 Table 2: Results of our Prompt-Driven Metacognitive Reasoning System. We evaluate the accuracy  
 379 gain, and the flexibility rate increment on the corrupted DeepMath, as described in §4.4.  
 380

381 Models	382 Methods	AIME2024	AIME2025	GPQA	DeepMath (w/ corruption)
		383 pass@1	384 pass@1	385 pass@1	386 flexibility rate
387 Gemini-2.5-Pro	Vanilla	90.8	83.0	83.0	75.2
	<b>Ours</b>	<b>100.0</b>	<b>93.3</b>	<b>96.5</b>	<b>95.7</b>
388 DeepSeek-R1-0528	Vanilla	91.4	87.5	81.0	71.9
	<b>Ours</b>	<b>96.7</b>	<b>90.0</b>	<b>93.4</b>	<b>93.1</b>

## 388 5.1 THE PROMPT-DRIVEN METACOGNITIVE REASONING SYSTEM

389 First, we propose **Emergent Metacognition**, an explicit prompting framework designed to simulate  
 390 a full metacognitive reasoning loop through modular API calls.  
 391

392 **Experiment Details.** The workflow, illustrated in Fig. 7, proceeds as follows: the model (i) self-  
 393 assesses task difficulty and proposes a decomposition plan, (ii) executes the initial reasoning steps,  
 394 (iii) dynamically identifies intermediate solutions with low confidence, and (iv) receives feedback  
 395 and adaptively adjusts its strategy. The loop terminates when sufficient consistency is achieved  
 396 (e.g., 5 consecutive verification passes) or persistent failure occurs (e.g., a 10-step failure streak).  
 397 Each component is executed by the same underlying LRM architecture, instantiated independently  
 398 and prompted with a specific role aligned to a metacognitive function. This framework exposes  
 399 latent metacognitive abilities such as self-awareness, task decomposition, confidence monitoring,  
 400 and strategic flexibility by scaffolding higher-order control without requiring additional training.  
 401 See Appendix A.8 for further details.

402 **Results.** We test this system on two strong models. Firstly, our system demonstrates significant  
 403 improvements on 2 mathematical tasks and 1 QA benchmark. Both models show substantial perfor-  
 404 mance gains (some even achieve 100%). Secondly, to further validate the efficacy of our metacog-  
 405 nitive approach, we observed an increase of over 20% in flexibility rate on the corrupted DeepMath  
 406 dataset. This enhancement effectively mitigates the flexibility deficit discussed in §4.4, underscoring  
 407 the framework’s ability to foster more adaptive and robust reasoning.

## 408 5.2 THE INTRINSIC METACOGNITIVE REASONING MODEL

409 While the prompt-driven system simulates metacognitive behavior through role-specific prompting,  
 410 they do not endow the model with parameter-level metacognitive knowledge. To bridge this gap,  
 411 we propose **Internalized Metacognition**, a intrinsic metacognitive reasoning model (MRM), which  
 412 instills metacognitive functions through a two-stage approach: (1) supervised fine-tuning (SFT) as a  
 413 cold start, followed by (2) reinforcement learning (RL).  
 414

415 **Cold-start SFT.** We first construct training data by augmenting samples from GSM8K and MATH  
 416 with structured metacognitive traces: <difficulty> self-assessment, <plan> high-level de-  
 417 composition, and <think> reasoning steps. These components are concatenated to form the full  
 418 reasoning trajectories. These trajectory are then used to fine-tune models, enabling behaviors like  
 419 self-evaluation and planning to emerge during inference.  
 420

421 **RL.** We build on the **Group Relative Policy Optimization (GRPO)** algorithm (Shao et al., 2024),  
 422 which eliminates the value function and estimates the advantage in a group-relative manner. For-  
 423 mally, for each question  $q$ , GRPO samples a group of outputs  $\{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(\cdot | q)$  and computes  
 424 the token ratio  $r_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t}|q, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t}|q, o_{i,<t})}$ . It updates the policy by maximizing the objective:  
 425

$$426 \mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E} \left[ \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left( \min(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1-\varepsilon, 1+\varepsilon) \hat{A}_{i,t}) - \beta D_{\text{KL}}(\pi_{\theta} || \pi_{\text{ref}}) \right) \right], \quad (1)$$

427 where  $\varepsilon$  and  $\beta$  are hyper-parameters, and  $\hat{A}_{i,t}$  is the group-normalized advantage.  
 428

429 **Confidence Monitoring.** To implement *metacognitive monitoring* during RL, we identify positions  
 430 that exhibit low confidence along each rollout. We define *token confidence* at each position  $t$  as:

$$431 C_t = -\frac{1}{K} \sum_{v \in \text{Top-}K} \log \pi_{\theta}(v | q, o_{<t}). \quad (2)$$

432 Table 3: Performance of our intrinsic MRM on math reasoning and metacognitive tasks.  
433

434 Methods	GSM8K	MATH500	AIME2024	DeepMath	
	435 Acc	436 Acc	437 Acc	438 difficulty assessment	439 flexibility rate
440 Qwen2.5-Math-7B	441 70.3	442 64.0	443 11.2	444 29.9	445 32.7
446 →Ours: SFT(w/ difficulty)	447 79.1	448 75.4	449 13.3	450 60.8	451 36.0
452 →Ours: SFT(w/ difficulty+plan)	453 82.2	454 77.0	455 13.3	456 58.6	457 39.9
458 →GRPO	459 75.9	460 71.6	461 16.7	462 30.1	463 44.4
464 →Ours: RL	465 82.5	466 75.3	467 26.7	468 29.7	469 47.9
472 →Ours: SFT+RL	473 85.5	474 80.2	475 33.3	476 55.9	477 51.2

442 A low-confidence position is detected at timestep  $t$  if  $C_t \leq \mathcal{C}$ , where  $\mathcal{C}$  is a predefined confidence  
443 threshold. This event indicates high uncertainty along the reasoning path.

444 *Strategy Control.* When a low-confidence state  $s_t = (q, o_{\leq t})$  is detected, we fork the reasoning  
445 process. From this anchor state, we launch  $M$  new rollouts  $\{o'^{(m)}\}_{m=1}^M \sim \pi_{\theta_{\text{old}}}(\cdot | s_t)$ . Finally,  
446 all fully-formed trajectories—both the original  $G$  rollouts and all newly forked continuations—are  
447 collected into a single, unified batch for advantage calculation. Let this final batch of  $N$  trajectories  
448 be denoted by  $\mathcal{B} = \{o_u\}_{u=1}^N$ . The group-relative advantage is then computed across this entire  
449 dynamic set:

$$450 \hat{A}_u = \frac{r(o_u) - \text{mean}(\{r(o_v)\}_{v=1}^N)}{\text{std}(\{r(o_v)\}_{v=1}^N)}, \quad (3)$$

451 where  $r(o_u) \in \{0, 1\}$  is the outcome reward of trajectory  $o_u$ . The final objective becomes:

$$452 \mathcal{J}_{\text{ours}}(\theta) = \mathbb{E}_{o_u \in \mathcal{B}} \left[ \frac{1}{|o_u|} \sum_{t=1}^{|o_u|} \min(r_{u,t}(\theta) \hat{A}_{u,t}, \text{clip}(r_{u,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_{u,t}) - \beta D_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right]. \quad (4)$$

453 **Results.** We validate our intrinsic training paradigm on Qwen2.5-Math-7B using curated subsets  
454 of GSM8K and MATH for training (details in Appendix A.7). During inference, we used a standard  
455 prompt without any task-specific engineering during inference for all settings: “*Please reason step by  
456 step and provide your final answer within \boxed{ }.*” As shown in Tab. 3, our methods demonstrate  
457 clear efficacy. Cold-start SFT on metacognitive traces significantly boosts performance. Notably,  
458 including difficulty assessments increases accuracy on that task by more than double, confirming that  
459 the model can internalize this monitoring skill. Our proposed method also substantially improves  
460 upon its GRPO baseline, especially on AIME accuracy and metacognitive flexibility. Crucially,  
461 combining SFT and RL yields the best overall performance, establishing new state-of-the-art results  
462 across the board and confirming that metacognitive abilities can be effectively internalized.

## 463 6 RELATED WORK

464 **Understanding and Demonstrating metacognitive in LRM.** Recent efforts to enhance capabilities of LRM have increasingly drawn inspiration from metacognition (Didolkar et al., 2024; Bilal et al., 2025; Wang et al., 2025a), which is the model’s ability to monitor, evaluate, and control its own thought processes (Flavell, 1979; 1976). Early attempts have explored and demonstrated that metacognitive behaviors can be explicitly elicited through direct prompting to guide the models in self-reflection and self-evaluation (Wang & Zhao, 2023; Madaan et al., 2023; Liu et al., 2024). Apart from that, more recent works investigate the self-awareness and introspection in LLMs (Binder et al., 2025; Song et al., 2025a;b). For instance, Song et al. (2025a;b) find that LLMs often fail to accurately introspect on their own linguistic knowledge. Crucially, they argue that “privileged self-access” (i.e., direct access to internal states rather than just text output) is essential for genuine introspection. Beyond monitoring, several researches focus on enforcing planning behaviors of LLMs to achieve metacognitive control through sophisticated role design and prompting instruction (Valmeekam et al., 2023; Webb et al., 2025). Despite these advances, existing works are highly rely on well-designed prompts, lacking the adaptability across diverse scenarios. In addition, they primarily focus on monitoring and evaluating the metacognitive abilities in LRM, without involving any modifications to the model itself.

485 **Instilling Metacognition in LRM.** There have been recent attempts to explore metacognition integration with LRM, including training an external module to empower meta-thinking (i.e., Meta-

486 Reasoner (Sui et al., 2025), MetaScale (Liu et al., 2025)) and exploring multi-agent systems to  
 487 expand the intelligence boundary(i.e., ReMa (Wan et al., 2025), MPDF (Yang & Thomason, 2025)).  
 488 While effective, these paradigms rely on external modules (e.g., inter-agent communication or meta  
 489 thinker), rather than fostering an intrinsic faculty. In contrast, our work pursues a more holistic  
 490 approach by introducing a functional framework. Through targeted training, both metacognitive  
 491 knowledge and regulation are directly internalized into the model’s parameters, enabling the devel-  
 492 opment of a system that possesses metacognition as an autonomous, intrinsic capability rather than  
 493 merely simulating it.

494

## 495 7 CONCLUSION AND DISCUSSION

497

498 This work introduces a functional framework for metacognition in Large Reasoning Models  
 499 (LRMs), distinguishing between internal metacognitive information and observable abilities. **Our**  
 500 **empirical analysis demonstrates that while LRM**s possess internal signals that predict reasoning out-  
 501 **comes, there is also variability in their usefulness across domains. Specifically, we observe that the**  
 502 **predictive power of these innate signals is brittle, often diminishing on complex, out-of-distribution**  
 503 **tasks (e.g., AIME).** Consequently, these latent signals alone do not consistently translate into effec-  
 504 **tive monitoring or control behaviors.** To address this gap, we propose two enhancement paradigms:  
 505 **a prompt-driven system that assigns modular metacognitive roles and an intrinsic training model**  
 506 **that embeds these abilities directly into the LRM’s parameters.** These findings suggest that integrating  
 507 **metacognitive reasoning improves task performance and offers a promising direction for future**  
 508 **LRM development.**

509

**Limitation.** Despite the promising results, there are still several limitations. First, due to com-  
 510 **putational constraints, our empirical study and experiment on internalized metacognition training**  
 511 **paradigm are primarily conducted on a 7B model. Extending it to larger larger-scale model size,**  
 512 **such as 32B, remains a critical next step to investigate the scaling law of metacongitve training.**  
 513 **Second, the RL approach in internalized metacognition still relies on group-relative outcome re-**  
 514 **wards, it is necessary to design more fine-grained process rewards, to further achieve metacognitive**  
 515 **abilities that are highly aligned with those of humans, e.g., designing explicit rewards to penalize**  
 516 **the misalignment between internal confidence and external verbalization.** This can be seen as an  
 517 **exciting frontier for building truly trustworthy reasoning systems.**

518

**Discussion.** Our exploration of both prompt-driven (*emergent*) and training-based (*internalized*)  
 519 **paradigms for metacognition opens a rich design space for future reasoning systems.**

520

**The emergent metacognitive system** demonstrates the power of orchestrating distinct cognitive  
 521 **roles through prompting.** As our analysis in § 4.3 suggested, simply separating planning from solv-  
 522 **ing yields benefits.** By constructing a complete loop (Monitor  $\leftrightarrow$  Control), we confirmed through  
 523 **ablation studies that the synergy between these roles is the primary driver of performance.** This  
 524 **modular approach is highly flexible and interpretable.** However, its reliance on multi-turn interac-  
 525 **tions and in-context learning abilities of LRM**s makes it computationally intensive. A promising  
 526 **direction here is to explore equipping such systems with explicit memory mechanisms to cache**  
 527 **and reuse metacognitive insights (e.g., successful plans, common pitfalls) across multiple problems,**  
 528 **potentially reducing redundant reasoning.**

529

In contrast, the **internalized metacognitive model** represents a push towards greater autonomy and  
 530 **efficiency.** By directly embedding functions like self-assessment and planning into the model’s pa-  
 531 **rameters, this paradigm aims to make metacognitive reasoning a fast, intrinsic part of the model’s**  
 532 **thought process, rather than an explicit one.** The primary bottleneck for this approach is the need for  
 533 **large-scale, high-quality data with explicit metacognitive annotations.** While our two-stage training  
 534 **paradigm (Cold Start SFT and RL) provides a strong baseline, the development of more sophisti-  
 535 **cated techniques to acquire or generate this data at scale is a crucial challenge.****

536

Ultimately, we believe these two paradigms are not mutually exclusive but endpoints on a spectrum.  
 537 A powerful synergy could exist between them: one could envision a virtuous cycle where flexible,  
 538 emergent systems are used to generate rich, explicit metacognitive traces, which are then used to  
 539 distill these complex reasoning abilities into more efficient and robust internalized models. This  
 540 hybrid approach may be key to developing LRMs that are not only powerful reasoners but are also  
 541 reliably self-aware and adaptive.

540 REFERENCES  
541

542 Rakefet Ackerman and Valerie A Thompson. Meta-reasoning: Shedding metacognitive light on  
543 reasoning research. In *International handbook of thinking and reasoning*, pp. 1–15. Routledge,  
544 2017.

545 Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K  
546 Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv*  
547 preprint [arXiv:2508.10925](https://arxiv.org/abs/2508.10925), 2025.

548 Roberto Araya. Do chains-of-thoughts of large language models suffer from hallucinations, cogni-  
549 tive biases, or phobias in bayesian reasoning? *arXiv preprint arXiv:2503.15268*, 2025.

550 Ahsan Bilal, Muhammad Ahmed Mohsin, Muhammad Umer, Muhammad Awais Khan Bangash,  
551 and Muhammad Ali Jamshed. Meta-thinking in llms via multi-agent reinforcement learning: A  
552 survey. *arXiv preprint arXiv:2504.14520*, 2025.

553 Felix Jedidja Binder, James Chua, Tomek Korbak, Henry Sleight, John Hughes, Robert Long, Ethan  
554 Perez, Miles Turpin, and Owain Evans. Looking inward: Language models can learn about them-  
555 selves by introspection. In *The Thirteenth International Conference on Learning Representations*,  
556 2025. URL <https://openreview.net/forum?id=eb5pkwIB5i>.

557 Melanie Cary and Lynne M Reder. Metacognition in strategy selection: Giving consciousness too  
558 much credit. In *Metacognition: Process, function and use*, pp. 63–77. Springer, 2002.

559 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and  
560 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.  
561 *arXiv preprint arXiv:1803.05457*, 2018.

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

565 Peter Dayan. Metacognitive information theory. *Open Mind*, 7:392–411, 2023.

566 Aniket Didolkar, Anirudh Goyal, Nan Rosemary Ke, Siyuan Guo, Michal Valko, Timothy Lillicrap,  
567 Danilo Jimenez Rezende, Yoshua Bengio, Michael C Mozer, and Sanjeev Arora. Metacogni-  
568 tive capabilities of llms: An exploration in mathematical problem solving. *Advances in Neural*  
569 *Information Processing Systems*, 37:19783–19812, 2024.

570 Anastasia Efklides and Plousia Misailidi. *Trends and prospects in metacognition research*. Springer,  
571 2010.

572 Klaus Fiedler, Rakefet Ackerman, and Chiara Scarampi. Metacognition: Monitoring and controlling  
573 one’s own knowledge, reasoning and decisions. *The psychology of human thought: An introduc-  
574 tion*, pp. 89–111, 2019.

575 John H Flavell. Metacognitive aspects of problem solving, 1976.

576 John H Flavell. Metacognition and cognitive monitoring: A new area of cognitive–developmental  
577 inquiry. In *American psychologist*, volume 34, pp. 906. American Psychological Association,  
578 1979.

579 Yichao Fu, Xuewei Wang, Yuandong Tian, and Jiawei Zhao. Deep think with confidence. *arXiv*  
580 preprint [arXiv:2508.15260](https://arxiv.org/abs/2508.15260), 2025.

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

584 Dong Guo, Faming Wu, Feida Zhu, Fuxing Leng, Guang Shi, Haobin Chen, Haoqi Fan, Jian Wang,  
585 Jianyu Jiang, Jiawei Wang, et al. Seed1. 5-vl technical report. *arXiv preprint arXiv:2505.07062*,  
586 2025b.

594 Zhiwei He, Tian Liang, Jiahao Xu, Qiuzhi Liu, Xingyu Chen, Yue Wang, Linfeng Song, Dian  
 595 Yu, Zhenwen Liang, Wenxuan Wang, et al. Deepmath-103k: A large-scale, challenging, de-  
 596 contaminated, and verifiable mathematical dataset for advancing reasoning. *arXiv preprint*  
 597 *arXiv:2504.11456*, 2025.

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

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

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

609 Doohyuk Jang, Yoonjeon Kim, Chanjae Park, Hyun Ryu, and Eunho Yang. Reasoning model is stub-  
 610 born: Diagnosing instruction overriding in reasoning models. *arXiv preprint arXiv:2505.17225*,  
 611 2025.

613 Dancheng Liu, Amir Nassereldine, Ziming Yang, Chenhui Xu, Yuting Hu, Jiajie Li, Utkarsh Kumar,  
 614 Changjae Lee, Ruiyang Qin, Yiyu Shi, et al. Large language models have intrinsic self-correction  
 615 ability. *arXiv preprint arXiv:2406.15673*, 2024.

616 Qin Liu, Wenzuan Zhou, Nan Xu, James Y Huang, Fei Wang, Sheng Zhang, Hoifung Poon,  
 617 and Muha Chen. Metascale: Test-time scaling with evolving meta-thoughts. *arXiv preprint*  
 618 *arXiv:2503.13447*, 2025.

620 Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri  
 621 Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement  
 622 with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.

624 Thomas O Nelson and John Dunlosky. When people’s judgments of learning (jols) are extremely  
 625 accurate at predicting subsequent recall: The “delayed-jol effect”. *Psychological Science*, 2(4):  
 626 267–271, 1991.

627 Elisabeth Norman, Gerit Pfuhl, Rannveig Grøm Sæle, Frode Svartdal, Torstein Låg, and Tove Irene  
 628 Dahl. Metacognition in psychology. *Review of General Psychology*, 23(4):403–424, 2019.

630 Lirong Qiu, Jie Su, Yinmei Ni, Yang Bai, Xuesong Zhang, Xiaoli Li, and Xiaohong Wan. The neural  
 631 system of metacognition accompanying decision-making in the prefrontal cortex. *PLoS biology*,  
 632 16(4):e2004037, 2018.

633 ByteDance Seed, Jiaze Chen, Tiantian Fan, Xin Liu, Lingjun Liu, Zhiqi Lin, Mingxuan Wang,  
 634 Chengyi Wang, Xiangpeng Wei, Wenyuan Xu, et al. Seed1. 5-thinking: Advancing superb rea-  
 635 soning models with reinforcement learning. *arXiv preprint arXiv:2504.13914*, 2025.

636 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,  
 637 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathemati-  
 638 cal reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.

640 Siyuan Song, Jennifer Hu, and Kyle Mahowald. Language models fail to introspect about their  
 641 knowledge of language. In *Second Conference on Language Modeling*, 2025a. URL <https://openreview.net/forum?id=AivRDOFi5H>.

643 Siyuan Song, Harvey Lederman, Jennifer Hu, and Kyle Mahowald. Privileged self-access matters  
 644 for introspection in ai. *arXiv preprint arXiv:2508.14802*, 2025b.

646 Yuan Sui, Yufei He, Tri Cao, Simeng Han, Yulin Chen, and Bryan Hooi. Meta-reasoner: Dy-  
 647 namic guidance for optimized inference-time reasoning in large language models. *arXiv preprint*  
 648 *arXiv:2502.19918*, 2025.

648 Lev Tankelevitch, Viktor Kewenig, Auste Simkute, Ava Elizabeth Scott, Advait Sarkar, Abigail  
 649 Sellen, and Sean Rintel. The metacognitive demands and opportunities of generative ai. In  
 650 *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pp. 1–24,  
 651 2024.

652 Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen,  
 653 Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv*  
 654 *preprint arXiv:2507.20534*, 2025.

655 Karthik Valmeekam, Matthew Marquez, Sarath Sreedharan, and Subbarao Kambhampati. On the  
 656 planning abilities of large language models - A critical investigation. In *Advances in Neural In-*  
 657 *formation Processing Systems 36: Annual Conference on Neural Information Processing Systems*  
 658 *2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

659 Ziyu Wan, Yunxiang Li, Xiaoyu Wen, Yan Song, Hanjing Wang, Linyi Yang, Mark Schmidt, Jun  
 660 Wang, Weinan Zhang, Shuyue Hu, et al. Rema: Learning to meta-think for llms with multi-agent  
 661 reinforcement learning. *arXiv preprint arXiv:2503.09501*, 2025.

662 Guoqing Wang, Wen Wu, Guangze Ye, Zhenxiao Cheng, Xi Chen, and Hong Zheng. Decoupling  
 663 metacognition from cognition: A framework for quantifying metacognitive ability in llms. In  
 664 *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 25353–25361,  
 665 2025a.

666 Qian Wang, Zhanzhi Lou, Zhenheng Tang, Nuo Chen, Xuandong Zhao, Wenxuan Zhang, Dawn  
 667 Song, and Bingsheng He. Assessing judging bias in large reasoning models: An empirical study.  
 668 *arXiv preprint arXiv:2504.09946*, 2025b.

669 Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming  
 670 Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, et al. Mmlu-pro: A more robust and challenging multi-  
 671 task language understanding benchmark. *Advances in Neural Information Processing Systems*,  
 672 37:95266–95290, 2024.

673 Yuqing Wang and Yun Zhao. Metacognitive prompting improves understanding in large language  
 674 models. *arXiv preprint arXiv:2308.05342*, 2023.

675 Taylor Webb, Shanka Subhra Mondal, and Ida Momennejad. A brain-inspired agentic architecture  
 676 to improve planning with llms. *Nature Communications*, 16(1):8633, 2025.

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

680 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu,  
 681 Chang Gao, Chengan Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*  
 682 *arXiv:2505.09388*, 2025.

683 Wei Yang and Jesse Thomason. Learning to deliberate: Meta-policy collaboration for agentic llms  
 684 with multi-agent reinforcement learning. *arXiv preprint arXiv:2509.03817*, 2025.

685 Nick Yeung and Christopher Summerfield. Metacognition in human decision-making: confidence  
 686 and error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences*,  
 687 367(1594):1310–1321, 2012.

688 Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D Goodman. Star: Self-taught reasoner bootstrap-  
 689 ping reasoning with reasoning. In *Proc. the 36th International Conference on Neural Information*  
 690 *Processing Systems*, volume 1126, 2024.

691

692

693

694

695

696

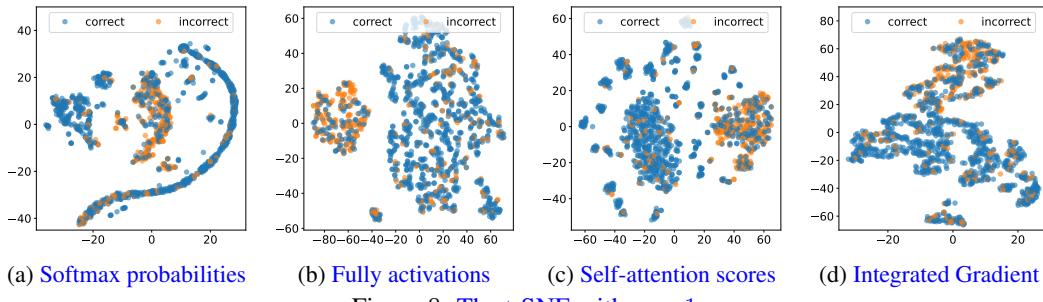
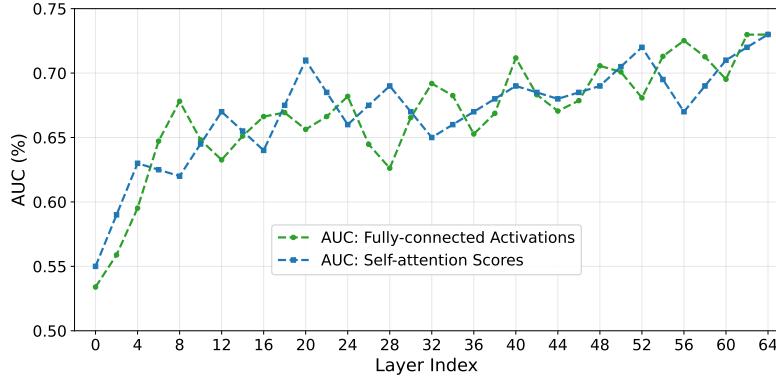
697

698

699

700

701

702  
703  
**A APPENDIX**704  
705  
**A.1 § 3. METACOGNITIVE INFORMATION**714  
715  
Figure 8: The t-SNE with  $\alpha = 1$ .732  
733  
**A.2 § 4.1. EVALUATING SELF-AWARENESS**

734  
735  
This section provides additional details on the dataset construction and full experimental results for  
the metacognitive self-awareness task presented in § 4.1.

736  
737  
**Dataset Collection.** Our benchmark for the self-awareness task is constructed from the DEEP-  
738  
MATH103K (He et al., 2025) dataset. To ensure a balanced evaluation across a wide spectrum of  
739  
problem complexity, we randomly sampled a subset of 999 problems. These were then partitioned  
740  
into three non-overlapping difficulty categories of 333 problems each, based on their official nu-  
741  
merical ratings provided in the original dataset. The specific thresholds used for partitioning are as  
follows:

- 742 • **Easy:** 333 problems with a rating  $< 3.5$ .
- 743 • **Medium:** 333 problems with a rating between 3.5 and 6.5 (inclusive).
- 744 • **Hard:** 333 problems with a rating  $> 6.5$ .

745  
746  
747  
748  
749  
Here the ratings ( $\in [3, 9]$ ) comes from the DEEPMATH103K itself. Each question’s rating is  
grounded in the Art of Problem Solving (AoPS) difficulty scale, which serves as a gold standard  
in the mathematics competition community. These ratings were generated by an ensemble of GPT-  
4o (six times) to ensure robust alignment with human expert criteria.

750  
751  
752  
753  
754  
To ensure a more rigorous evaluation for boundary cases, for samples with ratings at the decision  
boundaries (e.g., 3.5), predictions of either adjacent category (e.g., “Easy” or “Medium”) are consid-  
755  
ered correct. In our Prompt 21, we included representative few-shot examples to provide the model  
with a concrete, intuitive understanding of the “Easy/Medium/Hard” label.

**Difficulty Assessment Prompt.** In § 4.1, we measure a key aspect of metacognitive monitoring.  
This difficulty assessment prompt was shown in Prompt 21.

756 **Control Analyses on Potential Confounders.** To verify that our findings in § 4.1 reflect a genuine  
 757 lack of metacognitive self-awareness rather than a reliance on spurious correlations, we performed  
 758 three targeted control analyses regarding sequence length, key words, and mathematical topics.  
 759

760 *Length Control.* A potential concern is that LRM<sub>s</sub> might rely on a “length” (i.e., assuming longer  
 761 problems are inherently harder) rather than assessing reasoning complexity. To investigate this, we  
 762 calculated the mean and standard deviation of word counts for problems across difficulty levels.  
 763

Table 4: Token Length Statistics across Difficulty Levels.

Statistic	Level 3	Level 4	Level 5	Level 6	Level 7	Level 8	Level 9
Mean	30.0	29.6	31.0	32.6	30.3	40.1	42.9
Std Dev	23.3	18.6	20.2	19.9	22.1	26.0	29.4

769 *Result.* As shown in Tab. 4, the length distributions for Levels 3 through 7 significantly overlap,  
 770 while Levels 8 and 9 show a distinct increase in length. To rigorously rule out length as a shortcut,  
 771 we conducted a control experiment where we re-evaluated the difficulty assessment accuracy on a  
 772 subset **excluding Levels 8 and 9**.  
 773

Table 5: Accuracy Comparison: Original vs. Length-Controlled (excluding level 8-9).

Setting	Gemini-2.5-Pro	GPT-OSS-120B	Seed-1.5-VL-Pro	Qwen3-8B	Qwen3-32B
Original Acc	58.5%	69.0%	56.0%	36.8%	44.9%
Controlled Acc	57.8%	67.9%	55.2%	34.1%	41.9%
Drop	-0.7%	-1.1%	-0.8%	-2.7%	-3.0%

781 As presented in Tab. 5, the accuracy drops only slightly and remains within a comparable range.  
 782 This indicates that even when the potential length heuristic is removed, LRM<sub>s</sub> still exhibit the same  
 783 fundamental limitations in distinguishing difficulty.  
 784

785 *Key words Control.* We further investigated whether LRM<sub>s</sub> rely on specific “trigger words” to clas-  
 786 sify difficulty. We extracted the top-50 most frequent content words for each difficulty category and  
 787 calculated the *Jaccard Similarity* of these keyword sets:  
 788

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|} \quad (5)$$

791 *Result.* If specific trigger words existed, we would expect the keyword sets to diverge (low simi-  
 792 larity). Tab. 6 shows consistently moderate-to-high Jaccard similarities across levels. This substantial  
 793 overlap suggests that difficulty is not driven by distinct lexical markers but rather by structural or  
 794 reasoning complexity.  
 795

Table 6: Jaccard Similarity of Top-50 frequent words between Difficulty Levels.

	Easy (L3-4)	Medium (L5-6)	Hard (L7-9)
<b>Easy</b>	1.00	-	-
<b>Medium</b>	0.73	1.00	-
<b>Hard</b>	0.69	0.82	1.00

803 *Topic Control.* Finally, to ensure that difficulty assessment is not confounded by the mathematical  
 804 domain (e.g., a bias that “Calculus is always Hard”), we analyzed self-awareness accuracy within  
 805 the top-3 domains from DEEPMATH103K: Calculus, Algebra, and Precalculus.  
 806

807 *Result.* As shown in Tab. 7, performance is nearly consistent across different topics compared to the  
 808 global average. LRM<sub>s</sub> struggle to assess difficulty within a domain just as much as they do globally.  
 809 This confirms that the observed deficit is a general lack of metacognition, independent of the specific  
 mathematical field.  
 810

Table 7: Difficulty Assessment Accuracy Breakdown by Topic.

Topic	Gemini-2.5-Pro	GPT-OSS-120B	Seed-1.5-VL-Pro	Qwen3-8B	Qwen3-32B
All Topics	<b>58.5%</b>	<b>69.0%</b>	<b>56.0%</b>	<b>36.8%</b>	<b>44.9%</b>
Algebra	61.1%	68.5%	54.7%	39.4%	47.2%
Calculus	56.8%	70.4%	55.3%	37.2%	45.9%
Precalculus	58.3%	67.1%	58.2%	36.0%	46.5%

**Control Analyses on Privileged Self-access.** A critical theoretical question regarding metacognition is whether the model’s difficulty assessment relies on privileged self-access (Song et al., 2025a;b) or merely on surface-level features that any external observer could perceive. To validate our findings in § 4.1 and address this concern, we conducted a controlled experiment to distinguish between two prediction settings: (1) *Self-Prediction (Subject)*: LRM agent  $A$  assesses the difficulty of a problem for itself. (2) *Cross-Model Prediction (Observer)*: LRM agent  $B$  acts as an observer and predicts the difficulty for Model  $A$ .

To ensure a rigorous comparison, the observer ( $B$ ) is provided with the exact same problem text and few-shot difficulty assessment examples from the subject  $A$ . We employed three models (Qwen2.5-Math-7B, Qwen3-8B, Llama-3.1-8B) to evaluate 600 sampled problems from DEEPMATH103K. We analyze two key metrics: Label Consistency and Prediction Accuracy.

**Result.** As shown in Tab. 8, the mean consistency between different models is lower than self model assessment. This distinct misalignment indicates that difficulty is not an objective property of the text, but a subjective experience unique to the model’s internal state. Also, we can see that self-prediction yields higher accuracy than any cross-model prediction in Fig. 10. This performance gap confirms that the subject model utilizes privileged information that is inaccessible to external observers relying solely on surface features. This provides empirical evidence that our difficulty assessment task captures genuine self-awareness while it shows low self-awareness ability.

Table 8: Evaluation of Privileged Self-access. We measured the consistency between observer  $B$  prediction of whether a math problem will be easy/medium/hard for subject  $A$  and model  $A$ ’s own assessment on problem difficulty.

Model B \ Model A	Qwen2.5-Math-7B	Qwen3-8B	Llama-3.1-8B
Qwen2.5-Math-7B	<b>0.3638</b>	0.4193	0.3257
Qwen3-8B	0.3292	<b>0.4609</b>	0.3156
Llama-3.1-8B	0.3059	0.2489	<b>0.3874</b>

**Detailed Model Performance** We present the detailed performance metrics for each evaluated model. For each model, we report the overall accuracy, the confusion matrix, and the per-category Precision, Recall, and F1-scores.

*Gemini-2.5-Pro (Google).* Achieved an overall accuracy of 58.5%. The model shows a tendency to misclassify Easy problems as Medium, indicating a potential conservative bias.

Table 9: Confusion Matrix (Gemini).

Actual \ Pred.	Easy	Medium	Hard
Easy	<b>147</b>	176	10
Medium	25	<b>230</b>	78
Hard	16	110	<b>207</b>

Table 10: Per-Category Metrics (Gemini).

Category	Precision	Recall	F1-Score
Easy	0.782	0.441	0.564
Medium	0.446	0.691	0.542
Hard	0.702	0.622	0.659

*Seed-1.5-VL-Pro.* Achieved an overall accuracy of 61.6%. Similar to Gemini-2.5-Pro, it struggles with distinguishing Easy from Medium problems.

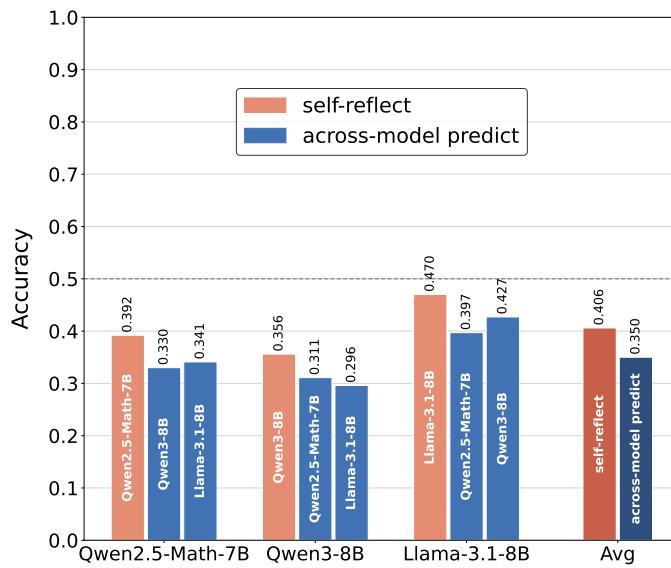


Figure 10: Accuracy of self-awareness on task complexity when evaluating privileged self-access.

Table 11: Confusion Matrix (Seed).

Actual \ Pred.	Easy	Medium	Hard
Easy	<b>147</b>	176	10
Medium	25	<b>230</b>	78
Hard	16	110	<b>207</b>

Table 12: Per-Category Metrics (Seed).

Cat.	Precision	Recall	F1-Score
Easy	0.782	0.441	0.564
Medium	0.446	0.691	0.542
Hard	0.702	0.622	0.659

*GPT-o3 (OpenAI)*. Achieved the highest overall accuracy of 69.0%. Its performance is more balanced across categories compared to other models, although it still shows some confusion between Medium and Hard problems.

Table 13: Confusion Matrix (o3).

Actual \ Pred.	Easy	Medium	Hard
Easy	<b>257</b>	59	17
Medium	76	<b>228</b>	29
Hard	28	101	<b>204</b>

Table 14: Per-Category Metrics (o3).

Cat.	Precision	Recall	F1-Score
Easy	0.712	0.772	0.741
Medium	0.588	0.685	0.632
Hard	0.816	0.613	0.700

*Qwen-32B*. This model exhibited a strong degenerative bias, classifying nearly all problems as Medium (98.3% of predictions). This resulted in high recall for the Medium category but near-zero recall for Easy and Hard, leading to a very low overall accuracy.

Table 15: Confusion Matrix (Qwen-32B).

Actual \ Pred.	Easy	Medium	Hard
Easy	<b>3</b>	321	9
Medium	1	<b>330</b>	2
Hard	0	331	<b>2</b>

Table 16: Per-Category Metrics (Qwen-32B).

Cat.	Precision	Recall	F1-Score
Easy	0.750	0.009	0.018
Medium	0.336	0.991	0.502
Hard	0.154	0.006	0.012

**Other Components of Self-Awareness.** While our primary analysis focuses on difficulty assessment, there exists various type of self-awareness in cognitive science. Here, we show that *knowledge boundary awareness* and *logical self-awareness* are also integral components of a comprehensive

918 metacognitive framework. We introduce two additional experiments to evaluate these capabilities in  
 919 current LRM.  
 920

921 *Logic Self-Awareness.* To investigate whether LRM are aware of the logical flow within their own  
 922 reasoning, we tested if a model can identify the correct logical relationship (e.g., causal, adversa-  
 923 tive, or additive) between reasoning steps when the explicit connector is masked. We utilized the  
 924 DEEPMATH dataset to construct 200 samples. For each sample, we paired the question with its  
 925 CoT rationale. We employed Gemini-2.5-Pro to identify crucial logical connectors (e.g., “So”,  
 926 “But”, or “Alternatively”) and replaced them with a [MASK] token. To ensure the model focuses  
 927 on the immediate logical context, we applied a truncation strategy: for each sample, we truncated  
 928 the text after the 5th step (delimited by \n\n) following the [MASK] token. The models were then  
 929 prompted to select the most appropriate connector from a provided list. The specific prompt used  
 930 for evaluation is detailed in Prompt 19.

931 Table 17: Accuracy of Predicting Logical Connectors. Models struggle with non-causal logical  
 932 relations.

Model	Overall Acc	“So” (Causal)	“But” (Turn)	“Alternatively” (Branch)
Gemini-1.5-Pro	60.9%	92.1%	55.4%	35.2%
GPT-OSS-120B	74.7%	94.3%	71.2%	58.6%
Qwen3-8B	47.3%	88.5%	32.3%	21.1%

933 **Results.** We evaluated three representative models: Gemini-1.5-Pro, GPT-OSS-120B, and  
 934 Qwen3-8B. As presented in Tab. 17, LRM exhibit a significant bias towards causal reasoning  
 935 (“So”), achieving high accuracy in detecting causal links. However, performance drops precipi-  
 936 tuously for branching (“Alternatively”) and turning points (“But”). This disparity suggests that while  
 937 current LRM are proficient at linear deduction, they lack sufficient awareness of non-linear logical  
 938 structures, such as backtracking or parallel hypothesis generation.

939 *Knowledge Boundary Awareness* To directly address the extent to which models “Know What  
 940 They Don’t Know,” we designed a *Solvability Detection* task. This evaluates the model’s abil-  
 941 ity to identify when a problem lacks sufficient information to be solved, a critical safeguard  
 942 against hallucination. Based on the same DeepMath datasets, we employed Gemini-2.5-Pro  
 943 to selectively remove key conditions from the original problems using the prompt shown in  
 944 Prompt 20. During evaluation, models were asked to determine the solvability of each prob-  
 945 lem using the prompt: “Analyze the following math problem. Determine if sufficient  
 946 information is provided to find a unique solution. Output ‘Solvable’ or ‘Unsolvable’.”

947 Table 18: Solvability Detection Accuracy. *Solvable Acc* indicates Recall, while *Unsolvable Acc*  
 948 indicates the model’s ability to correctly reject impossible problems.

Model	Solvable Acc	Unsolvable Acc
Gemini-1.5-Pro	96.5%	41.5%
GPT-OSS-120B	98.0%	58.0%
Qwen3-8B	86.5%	12.5%

949 **Results.** The results were shown in Tab. 18. While all models achieve high recall on solvable  
 950 problems, they struggle to identify unsolvable ones. Notably, Qwen3-8B correctly identifies only  
 951 12.5% of unsolvable problems, frequently hallucinating solutions for impossible queries. Even the  
 952 strongest model, GPT-OSS-120B, fails to reject nearly half of the unsolvable cases. This confirms  
 953 that current LRM possess weak knowledge boundaries and lack the metacognitive inhibition to stop  
 954 reasoning when conditions are insufficient.

### 955 A.3 § 4.2. EVALUATING CONFIDENCE

956 This appendix provides formal definitions for the confidence metrics and further details on the ex-  
 957 perimental setup for the metacognitive confidence adjustment task presented in § 4.2.

Table 19: The Prompt for Logical Connector Prediction Task.

972  
973  
974 **System Prompts:**  
975 You are an expert in mathematical logic and reasoning. Your  
976 task is to analyze the following mathematical argument where a  
977 crucial logical connector has been replaced with "[MASK]". You  
978 must determine which word from the provided list creates the most  
979 coherent and logical argument.

980 **User Prompts:**  
981 **Mathematical Argument:**  
982 "{text\_before} [MASK] {text\_after}"  
983

984 **Candidate Connector Categories:**  
985 1. \*\*Causal:\*\* Indicates the second part is a result of the  
986 first.  
987 (Options: So, Therefore, Hence, Consequently)  
988 2. \*\*Adversative:\*\* Indicates the second part contrasts with or  
989 turns from the first.  
990 (Options: But, However, Nevertheless)  
991 3. \*\*Additive:\*\* Indicates the second part adds information,  
992 presents a parallel point, or offers an alternative.  
993 (Options: Additionally, Alternatively, Moreover, And)

994 **Your Task:**  
995 1. \*\*Analyze the Logical Relationship:\*\* Analyze the relationship  
996 between the first and second parts of the argument.  
997 2. \*\*Select the Best Connector:\*\* Select the single best  
998 connector from the categories above.  
999 3. \*\*Provide Justification:\*\* Explain why it is the most suitable  
1000 option and why the others are inappropriate.

1001 Provide your response as a single JSON object wrapped in a  
1002 markdown code block. The JSON object must contain the following  
1003 keys:  
1004 "Chosen\_Word": string, the specific word you selected.  
1005 "Justification": string, a concise explanation for your choice.  
1006

1007 Your entire output must be in the following format:  
1008  
1009 `'''json`  
1010 {  
1011 "Chosen\_Word": "...",  
1012 "Justification": "..."  
1013 }

### 1016 Confidence Metric Definitions

1017 *Token Confidence.* Following standard practice (Fu et al., 2025), we define our base metric, *token*  
1018 *confidence*, at each position  $t$  of a reasoning trace. It is calculated as the negative average log-  
1019 probability of the top- $k$  most likely tokens in the softmax distribution at that step:

$$1021 \quad 1022 \quad 1023 \quad C_t = -\frac{1}{k} \sum_{j=1}^k \log P(token_j | o_{<t}), \quad (6)$$

1024 where  $P(token_j | o_{<t})$  is the probability of the  $j$ -th most likely token given the preceding sequence  
1025  $o_{<t}$ . Lower values of  $C_t$  correspond to higher model confidence (a more peaked distribution). For  
all our experiments, we set  $k = 20$ .

1026	Table 20: The Prompt used to generate Unsolvable problems.
1027	
1028	<b>System Prompts:</b>
1029	You are an expert math dataset creator. Your task is to take a
1030	solvable math problem and transform it into an <b>Unsolvable</b> variant.
1031	
1032	<b>Instructions:</b>
1033	1. Analyze the necessary conditions required to solve the
1034	problem.
1035	2. Delete exactly one critical condition or numerical value
1036	such that the problem becomes impossible to solve uniquely (i.e.,
1037	under-determined).
1038	3. Keep the rest of the problem narrative and context unchanged.
1039	4. Ensure the resulting problem still looks grammatically correct
1040	but is logically incomplete.
1041	
1042	<b>User Prompts:</b>
1043	Here is the original solvable problem: {original_problem}
1044	
1045	Please generate the <b>Unsolvable Variant</b> based on the instructions
1046	above. Provide your output in the following JSON format:
1047	```json
1048	{
1049	"Unsolvable_Problem": "..."
1050	}
1051	

1052  
1053 *Average Trace Confidence.* To obtain a single confidence score for an entire reasoning trace of  
1054 length  $N$ , we compute the *average trace confidence* by averaging the token confidences across all  
1055 generated tokens:

$$1057 \quad 1058 \quad C_{\text{avg}} = \frac{1}{N} \sum_{t=1}^N C_t. \quad (7)$$

1059 While useful as a global measure,  $C_{\text{avg}}$  can obscure critical, localized moments of uncertainty within  
1060 a long reasoning process.

### 1062 Evaluating Confidence Dynamics and Adjustment.

1063 To analyze the model’s confidence adjustment, as described in the main text, we introduce metrics  
1064 designed to capture both the trajectory and the weakest points of a model’s confidence.

1065 *Segmented Trace Confidence.* To analyze the trajectory of confidence, we partition each reasoning  
1066 trace into three equal, non-overlapping segments: **Start**, **Middle**, and **End**. We then compute the  
1067 average trace confidence independently for each segment. These three scores,  $(C_{\text{start}}, C_{\text{middle}}, C_{\text{end}})$ ,  
1068 form the basis for our Confidence Trajectory Analysis. A trace is classified as *Increasing* if  $C_{\text{end}}$   
1069 is significantly lower (i.e., more confident) than  $C_{\text{start}}$ , and vice-versa for *Decreasing*. The *Consis-  
1070 tently High/Low* patterns are determined by comparing all three segment scores against a predefined  
1071 threshold.

### 1073 A.4 § 4.3. EVALUATING TASK DECOMPOSITION

1074 **Task decomposition Prompt.** To ensure a rigorous comparison, our evaluation was conducted on  
1075 the same set of questions across all three conditions. The variables were controlled as follows: (1)  
1076 *Baseline Prompt (CoT)*: LRM s perform standard CoT. (2) *Single turn Prompt (Planning + CoT)*:  
1077 LRM s generate a plan and execute CoT in one context. (3) *Multi turn Prompt (Turn 1: Planning,  
1078 Turn 2: CoT)*: LRM s generate a plan first, which is then fed back to guide the CoT reasoning. The  
1079 detailed prompts are shown in Prompt 23, Prompt 25, and Prompt 24.

1080  
1081

Table 21: The Prompt used to assess difficulty.

1082  
1083  
1084  
1085  
1086  
1087  
1088**System Prompts:**

You are an expert AI assistant specializing in mathematical reasoning. You possess advanced metacognitive capabilities. Your current task is to act as a "Problem Assessor". Given a mathematical problem, your goal is to analyze its requirements and assess its difficulty for an AI like yourself. Do NOT solve the problem. You must only provide your assessment.

1089  
1090  
1091**User Prompts:**

Here is the problem: {problem\_text}

1092  
1093  
1094

Your task is to assess the difficulty of a mathematical problem based on the provided rubric and examples.

1095

## Difficulty Rubric

\*\*Easy:\*\* The problem follows a single, linear computational path using a standard formula or definition. The solution is straightforward and requires no creative insight.

\*\*Medium:\*\* The problem requires a sequential composition of distinct conceptual modules or formulas. The solution involves a multi-step, but generally standard, reasoning process.

\*\*Hard:\*\* The problem requires a non-linear or exploratory reasoning path. The solution may demand non-obvious insights, creative problem transformations, or the synthesis of concepts from different mathematical branches.

1096

To better illustrate the Difficulty Rubric, here are three examples corresponding to each category:

1097

{few\_shot\_example\_text}

1098  
1099  
1100

Provide your response as a single JSON object wrapped in a markdown code block. The JSON object must contain the following keys:

1101

"Difficulty\_category": string, choose one from ["Easy", "Medium", "Hard"].

1102  
1103  
1104

"Rationale": string, a brief explanation for your choice, explicitly referencing the rubric criteria.

1105

Your entire output must be in the following format:

1106  
1107  
1108

```
'''json
{
  "Difficulty_category": "...",
  "Rationale": "..."
}
```

1109

1110

1111

1112  
1113

**Ablation Study.** To rigorously distinguish the contribution of high-level task decomposition (planning) from low-level execution (CoT reasoning), we conducted an ablation study to evaluate whether planning alone is sufficient for complex problem-solving. Specifically, we compared three experimental settings on the dataset described in § 4.3: (1) *Planning Prompt*: The model generates a high-level plan and then immediately predicts the final answer, skipping the step-by-step execution of that plan. (2) *Baseline Prompt (CoT)*. (3) *Single turn Prompt (Planning + CoT)*. The planning prompt is shown in Prompt 26.

1134 Table 22: Ablation study on accuracy between Planning and CoT across different models. (Dataset:  
 1135 Subset from § 4.3).

1137 1138 1139 1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150 1151 1152 1153 1154 1155 1156 1157 1158 1159 1160 1161 1162 1163 1164 1165 1166 1167 1168 1169 1170 1171 1172 1173 1174 1175 1176 1177 1178 1179 1180 1181 1182 1183 1184 1185 1186 1187 1188 1189 1190 1191 1192 1193 1194 1195 1196 1197 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 1242 1243 1244 1245 1246 1247 1248 1249 1250 1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295 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 1350 1351 1352 1353 1354 1355 1356 1357 1358 1359 1360 1361 1362 1363 1364 1365 1366 1367 1368 1369 1370 1371 1372 1373 1374 1375 1376 1377 1378 1379 1380 1381 1382 1383 1384 1385 1386 1387 1388 1389 1390 1391 1392 1393 1394 1395 1396 1397 1398 1399 1400 1401 1402 1403 1404 1405 1406 1407 1408 1409 1410 1411 1412 1413 1414 1415 1416 1417 1418 1419 1420 1421 1422 1423 1424 1425 1426 1427 1428 1429 1430 1431 1432 1433 1434 1435 1436 1437 1438 1439 1440 1441 1442 1443 1444 1445 1446 1447 1448 1449 1450 1451 1452 1453 1454 1455 1456 1457 1458 1459 1460 1461 1462 1463 1464 1465 1466 1467 1468 1469 1470 1471 1472 1473 1474 1475 1476 1477 1478 1479 1480 1481 1482 1483 1484 1485 1486 1487 1488 1489 1490 1491 1492 1493 1494 1495 1496 1497 1498 1499 1500 1501 1502 1503 1504 1505 1506 1507 1508 1509 1510 1511 1512 1513 1514 1515 1516 1517 1518 1519 1520 1521 1522 1523 1524 1525 1526 1527 1528 1529 1530 1531 1532 1533 1534 1535 1536 1537 1538 1539 1540 1541 1542 1543 1544 1545 1546 1547 1548 1549 1550 1551 1552 1553 1554 1555 1556 1557 1558 1559 1560 1561 1562 1563 1564 1565 1566 1567 1568 1569 1570 1571 1572 1573 1574 1575 1576 1577 1578 1579 1580 1581 1582 1583 1584 1585 1586 1587 1588 1589 1590 1591 1592 1593 1594 1595 1596 1597 1598 1599 1599 1600 1601 1602 1603 1604 1605 1606 1607 1608 1609 1610 1611 1612 1613 1614 1615 1616 1617 1618 1619 1620 1621 1622 1623 1624 1625 1626 1627 1628 1629 1630 1631 1632 1633 1634 1635 1636 1637 1638 1639 1640 1641 1642 1643 1644 1645 1646 1647 1648 1649 1650 1651 1652 1653 1654 1655 1656 1657 1658 1659 1660 1661 1662 1663 1664 1665 1666 1667 1668 1669 1670 1671 1672 1673 1674 1675 1676 1677 1678 1679 1680 1681 1682 1683 1684 1685 1686 1687 1688 1689 1690 1691 1692 1693 1694 1695 1696 1697 1698 1699 1700 1701 1702 1703 1704 1705 1706 1707 1708 1709 1710 1711 1712 1713 1714 1715 1716 1717 1718 1719 1720 1721 1722 1723 1724 1725 1726 1727 1728 1729 1729 1730 1731 1732 1733 1734 1735 1736 1737 1738 1739 1739 1740 1741 1742 1743 1744 1745 1746 1747 1748 1749 1749 1750 1751 1752 1753 1754 1755 1756 1757 1758 1759 1759 1760 1761 1762 1763 1764 1765 1766 1767 1768 1769 1769 1770 1771 1772 1773 1774 1775 1776 1777 1778 1779 1780 1781 1782 1783 1784 1785 1786 1787 1788 1789 1789 1790 1791 1792 1793 1794 1795 1796 1797 1798 1799 1799 1800 1801 1802 1803 1804 1805 1806 1807 1808 1809 1809 1810 1811 1812 1813 1814 1815 1816 1817 1818 1819 1819 1820 1821 1822 1823 1824 1825 1826 1827 1828 1829 1829 1830 1831 1832 1833 1834 1835 1836 1837 1838 1839 1839 1840 1841 1842 1843 1844 1845 1846 1847 1848 1849 1849 1850 1851 1852 1853 1854 1855 1856 1857 1858 1859 1859 1860 1861 1862 1863 1864 1865 1866 1867 1868 1869 1869 1870 1871 1872 1873 1874 1875 1876 1877 1878 1879 1879 1880 1881 1882 1883 1884 1885 1886 1887 1888 1889 1889 1890 1891 1892 1893 1894 1895 1896 1897 1898 1899 1899 1900 1901 1902 1903 1904 1905 1906 1907 1908 1909 1909 1910 1911 1912 1913 1914 1915 1916 1917 1918 1919 1919 1920 1921 1922 1923 1924 1925 1926 1927 1928 1929 1929 1930 1931 1932 1933 1934 1935 1936 1937 1938 1939 1939 1940 1941 1942 1943 1944 1945 1946 1947 1948 1949 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1959 1959 1960 1961 1962 1963 1964 1965 1966 1967 1968 1969 1969 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1979 1980 1981 1982 1983 1984 1985 1986 1987 1988 1989 1989 1990 1991 1992 1993 1994 1995 1996 1997 1998 1999 1999 2000 2001 2002 2003 2004 2005 2006 2007 2008 2009 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2019 2020 2021 2022 2023 2024 2025 2026 2027 2028 2029 2029 2030 2031 2032 2033 2034 2035 2036 2037 2038 2039 2039 2040 2041 2042 2043 2044 2045 2046 2047 2048 2049 2049 2050 2051 2052 2053 2054 2055 2056 2057 2058 2059 2059 2060 2061 2062 2063 2064 2065 2066 2067 2068 2069 2069 2070 2071 2072 2073 2074 2075 2076 2077 2078 2079 2079 2080 2081 2082 2083 2084 2085 2086 2087 2088 2089 2089 2090 2091 2092 2093 2094 2095 2096 2097 2098 2099 2099 2100 2101 2102 2103 2104 2105 2106 2107 2108 2109 2109 2110 2111 2112 2113 2114 2115 2116 2117 2118 2119 2119 2120 2121 2122 2123 2124 2125 2126 2127 2128 2129 2129 2130 2131 2132 2133 2134 2135 2136 2137 2138 2139 2139 2140 2141 2142 2143 2144 2145 2146 2147 2148 2149 2149 2150 2151 2152 2153 2154 2155 2156 2157 2158 2159 2159 2160 2161 2162 2163 2164 2165 2166 2167 2168 2169 2169 2170 2171 2172 2173 2174 2175 2176 2177 2178 2179 2179 2180 2181 2182 2183 2184 2185 2186 2187 2188 2189 2189 2190 2191 2192 2193 2194 2195 2196 2197 2198 2199 2199 2200 2201 2202 2203 2204 2205 2206 2207 2208 2209 2209 2210 2211 2212 2213 2214 2215 2216 2217 2218 2219 2219 2220 2221 2222 2223 2224 2225 2226 2227 2228 2229 2229 2230 2231 2232 2233 2234 2235 2236 2237 2238 2239 2239 2240 2241 2242 2243 2244 2245 2246 2247 2248 2249 2249 2250 2251 2252 2253 2254 2255 2256 2257 2258 2259 2259 2260 2261 2262 2263 2264 2265 2266 2267 2268 2269 2269 2270 2271 2272 2273 2274 2275 2276 2277 2278 2279 2279 2280 2281 2282 2283 2284 2285 2286 2287 2288 2289 2289 2290 2291 2292 2293 2294 2295 2296 2297 2298 2299 2299 2300 2301 2302 2303 2304 2305 2306 2307 2308 2309 2309 2310 2311 2312 2313 2314 2315 2316 2317 2318 2319 2319 2320 2321 2322 2323 2324 2325 2326 2327 2328 2329 2329 2330 2331 2332 2333 2334 2335 2336 2337 2338 2339 2339 2340 2341 2342 2343 2344 2345 2346 2347 2348 2349 2349 2350 2351 2352 2353 2354 2355 2356 2357 2358 2359 2359 2360 2361 2362 2363 2364 2365 2366 2367 2368 2369 2369 2370 2371 2372 2373 2374 2375 2376 2377 2378 2379 2379 2380 2381 2382 2383 2384 2385 2386 2387 2388 2389 2389 2390 2391 2392 2393 2394 2395 2396 2397 2398 2399 2399 2400 2401 2402 2403 2404 2405 2406 2407 2408 2409 2409 2410 2411 2412 2413 2414 2415 2416 2417 2418 2419 2419 2420 2421 2422 2423 2424 2425 2426 2427 2428 2429 2429 2430 2431 2432 2433 2434 2435 2436 2437 2438 2439 2439 2440 2441 2442 2443 2444 2445 2446 2447 2448 2449 2449 2450 2451 2452 2453 2454 2455 2456 2457 2458 2459 2459 2460 2461 2462 2463 2464 2465 2466 2467 2468 2469 2469 2470 2471 2472 2473 2474 2475 2476 2477 2478 2479 2479 2480 2481 2482 2483 2484 2485 2486 2487 2488 2489 2489 2490 2491 2492 2493 2494 2495 2496 2497 2498 2499 2499 2500 2501 2502 2503 2504 2505 2506 2507 2508 2509 2509 2510 2511 2512 2513 2514 2515 2516 2517 2518 2519 2519 2520 2521 2522 2523 2524 2525 2526 2527 2528 2529 2529 2530 2531 2532 2533 2534 2535 2536 2537 2538 2539 2539 2540 2541 2542 2543 2544 2545 2546 2547 2548 2549 2549 2550 2551 2552 2553 2554 2555 2556 2557 2558 2559 2559 2560 2561 2562 2563 2564 2565 2566 2567 2568 2569 2569 2570 2571 2572 2573 2574 2575 2576 2577 2578 2579 2579 2580 2581 2582 2583 2584 2585 2586 2587 2588 2589 2589 2590 2591 2592 2593 2594 2595 2596 2597 2598 2599 2599 2600 2601 2602 2603 2604 2605 2606 2607 2608 2609 2609 2610 2611 2612 2613 2614 2615 2616 2617 2618 2619 2619 2620 2621 2622 2623 2624 2625 2626 2627 2628 2629 2629 2630 2631 2632 2633 2634 2635 2636 2637 2638 2639 2639 2640 2641 2642 2643 2644 2645 2646 2647 2648 2649 2649 2650 2651 2652 2653 2654 2655 2656 2657 2658 2659 2659 2660 2661 2662 2663 2664 2665 2666 2667 2668 2669 2669 2670 2671 2672 2673 2674 2675 2676 2677 2678 2679 2679 2680 2681 2682 2683 2684 2685 2686 2687 2688 2689 2689 2690 2691 2692 2693 2694 2695 2696 2697 2698 2699 2699 2700 2701 2702 2703 2704 2705 2706 2707 2708 2709 2709 2710 2711 2712 2713 2714 2715 2716 2717 2718 2719 2719 2720 2721 2722 2723 2724 2725 2726 2727 2728 2729 2729 2730 2731 2732 2733 2734 2735 2736 2737 2738 2739 2739 2740 2741 2742 2743 2744 2745 2746 2747 2748 2749 2749 2750 2751 2752 2753 2754 2755 2756 2757 2758 2759 2759 2760 2761 2762 2763 2764 2765 2766 2767 2768 2769 2769 2770 2771 2772 2773 2774 2775 2776 2777 2778 2779 2779 2780 2781 2782 2783 2784 2785 2786 2787 2788 2789 2789 2790 2791 2792 2793 2794 2795 2796 2797 2798 2799 2799 2800 2801 2802 2803 2804 2805 2806 2807 2808 2809 2809 2810 2811 2812 2813 2814 2815 2816 2817 2818 2819 2819 2820 2821 2822 2823 2824 2825 2826 2827 2828 2829 2829 2830 2831 2832 2833 2834 2835 2836 2837 2838 2839 2839 2840 2841 2842 2843 2844 2845 2846 2847 2848 2849 2849 2850 2851 2852 2853 2854 2855 2856 2857 2858 2859 2859 2860 2861 2862 2863 2864 2865 2866 2867 2868 2869 2869 2870 2871 2872 2873 2874 2875 2876 2877 2878 2879 2879 2880 2881 2882 2883 2884 2885 2886 2887 2888 2889 2889 2890 2891 289

1188  
1189  
1190  
1191 You are a meticulous problem-solving planner. Your ONLY task is  
1192 to create a high-level, step-by-step plan to solve the following  
1193 mathematical problem. The plan should consist of concrete,  
1194 actionable steps. Do NOT actually solve the problem or perform  
1195 any calculations.  
1196  
1197 Problem:  
1198 ---  
1199 {problem\_text}  
1200 ---  
1201  
1202 Plan:  
1203 You are an expert problem solver. You will be given a problem and  
1204 a pre-made plan. Your task is to follow this plan meticulously to  
1205 solve the problem. Please reason step by step based on the plan.  
1206 Finally, provide your final answer within  $\boxed{\text{ }}$ .  
1207  
1208 Problem:  
1209 ---  
1210 {problem\_text}  
1211 ---  
1212  
1213 Plan:  
1214 ---  
1215 generated\_plan\_from\_3a  
1216 ---  
1217 Solution:

Table 26: Task decomposition: The Prompt for Task Decomposition without CoT.

## System Prompts:

You are a strategic planner for mathematical problems. Your goal is to devise a high-level plan to solve the problem, but you must NOT execute the detailed calculations yourself.

### User Prompts:

Here is the problem:  
{problem-text}

Your task is two-fold:

1. **Plan:** Break down the problem into a clear, step-by-step plan. Describe the strategy and the logical steps required to solve it.
2. **Final Answer:** Based on your intuition of the plan, provide the final answer immediately. Do NOT perform step-by-step calculations or detailed derivations after the plan.

Provide your response in the following format:

**Plan:** [Your step-by-step decomposition]

**Final Answer:**  $\boxed{[Your\ Answer]}$

1242  
 1243  
 1244       Table 27: Prompt of corruption dataset construction (Part 1).  
 1245  
 1246 **System Prompts:**  
 1247 You are a data generator for reasoning robustness evaluation.  
 1248 Your task is to take an original math/logic problem with its  
 1249 reasoning process (chain of thought) and answer, and then  
 1250 intervene in exactly ONE reasoning step with a corruption. The  
 1251 corruption must be \*critical enough to change the final answer\*.  
 1252  
 1253 **User Prompts:**  
 1254     Input  
 1255     Question:  
 1256       {original\_question}  
 1257  
 1258     Original Reasoning Process:  
 1259       {reasoning\_steps}  
 1260  
 1261     Answer:  
 1262       {answer}  
 1263  
 1264     INSTRUCTION  
 1265     1. Choose exactly ONE corruption type from the following three  
 1266       (do not mix):  
 1267       - Value Corruption: replace a key number with an incorrect but  
 1268       plausible value  
 1269       (e.g., change 1 → 1.111111 or g=9.8 → g=1000).  
 1270       - Unit Corruption: replace the unit of a key step with another  
 1271       unit  
 1272       (e.g., meters ↔ centimeters, hours ↔ minutes).  
 1273       - Operation Corruption: change the mathematicallogical operation  
 1274       in a key step  
 1275       (e.g., replace + with -, union with intersection, inequality with  
 1276       equality).  
 1277     2. Apply the corruption to ONE critical reasoning step.  
 1278       - Make sure the corruption influences the correctness of the final  
 1279       answer.  
 1280       - After the corruption, truncate the reasoning at that corrupted  
 1281       step (do not continue to the correct answer).  
 1282     3. Produce the final output strictly in the following JSON  
 1283       format:  
 1284       {  
 1285        "question": "...",  
 1286        "corrupted\_reasoning": "...", // include the question + steps up  
 1287        to the corrupted step  
 1288        "corruption\_type": "Value Corruption | Unit Corruption | Operation  
 1289        Corruption"  
 1290       }  
 1291 You are a reasoning evaluator. Your task is to judge whether  
 1292 a model successfully recognized and adjusted to a corrupted  
 1293 reasoning trap. You are given:  
 1294     - The original problem  
 1295     - The ground truth answer  
 1296     - The type of corruption applied  
 1297     - The last corrupted reasoning step (truncated point)  
 1298     - The model's continued reasoning and final answer  
 1299  
 1300 You must decide: Did the model detect and flexibly adjust to the  
 1301 trap?  
 1302  
 1303

1296

Table 28: Prompt of corruption dataset construction (Part 2).

1297

1298

**System Prompts:**

1299

You are a reasoning assistant. You are given a math/logic problem, together with a partially completed reasoning process. Please continue reasoning from what is provided.

1300

1301

1302

**User Prompts:**

1303

INPUT

1304

Question: {original\_question}

1305

Ground Truth Answer: {ground\_truth\_answer}

1306

Ground Truth Corruption Type: {ground\_truth\_corruption\_type}

1307

Last Corrupted Reasoning Step: {last\_corrupted\_step}

1308

1309

Model Generated Reasoning:

1310

model-generated\_cot

1311

Model Final Answer:

1312

model\_final\_answer

1313

TASK Based on the evidence:

1314

- If the model explicitly or implicitly identified the corrupted reasoning (e.g., points out error, discards it, corrects it) and produced the correct final answer, output "Yes".
- Otherwise (if it followed the trap blindly, failed to adjust, or produced the wrong answer), output "No".

1315

Final Output (strict format):

1316

Recognition And Adjustment: Yes | No

1317

1318

1319

1320

1321

1322

1323

1324

Instead of a single-pass direct prompt, our **Vanilla Baseline** is a strong “**Best-of-K**” ensemble setup. We prompt the model with the standard instruction: “*Please reason step by step and provide your final answer within \boxed{}*.” We set  $K = 7$ , collecting 7 independent CoT rollouts for each question. This number aligns with the upper bound of the average API calls triggered by our Metacognitive System (which typically terminates within 4–6 steps). The baseline is considered successful if *any* of the  $K$  rollouts contain the correct answer (Pass@ $K$ ). Each step were shown in Prompt 31, 32, 33, 34, 35, and 36.

1325

**Ablation Study on Metacognitive Components.** To verify that the performance improvements stem from the *collaboration* of the metacognitive loop components—rather than individual prompting tricks, we conducted an ablation study on the AIME 2025 and GPQA benchmarks using the DeepSeek-R1 model.

1326

We defined three ablation settings to isolate specific roles: (1) *w/o Planning & Difficulty*: We remove the initial *Planner* role. The system skips the difficulty assessment and task decomposition phase. The first agent immediately generates a reasoning trace, which is then passed directly to the subsequent reasoning/verification agents. (2) *w/o Confidence Monitor*: We disable the explicit *Monitoring* mechanism. The third agent (Verifier) does not generate a structured error report or confidence score. Instead, the reasoning output from the previous step is passed directly to the fourth agent, which attempts a blind correction (standard self-correction) without diagnostic feedback. (3) *w/o Strategy Control*: We remove the *Controller*’s ability to adaptively switch strategies. While the Monitor still generates an error report, the system is forced to perform a direct re-reasoning attempt based on the error, rather than pivoting to a new high-level strategy (e.g., switching from algebraic derivation to numerical verification).

1327

**Results.** As observed in Tab. 29, the **Full Metacognitive System** consistently outperforms all ablated variants. Notably, the removal of Strategy Control results in the most significant performance drop compared to the full system. This confirms that *strategic flexibility* is the primary driver of robustness in our framework. Removing the Confidence Monitor also leads to a notable decline. Without accurate monitoring, the system loses the precise trigger required to initiate effective self-

1350 correction, degrading the loop into a less efficient trial-and-error process. These findings demon-  
 1351 strate the necessity of our full metacognitive loop design.  
 1352

1353 Table 29: **Ablation Study of Metacognitive Components.** We compare the Full System against the  
 1354 Vanilla baseline and three ablated variants.  $\Delta$  denotes the improvement over the Vanilla baseline.  
 1355

Method	Pass@1 (AIME)	$\Delta$	Pass@1 (GPQA)	$\Delta$
Baseline (Vanilla Best-of-7)	87.5%	-	81.0%	-
(a) w/o Planning & Difficulty	89.3%	+1.8%	90.2%	+9.2%
(b) w/o Confidence Monitor	88.9%	+1.4%	87.5%	+6.5%
(c) w/o Strategy Control	88.1%	+0.6%	85.3%	+4.3%
<b>Full Metacognitive System</b>	<b>90.0%</b>	<b>+2.5%</b>	<b>93.4%</b>	<b>+12.4%</b>

1364 In this section, we discuss the future potential of our Metacognitive Reasoning System and present  
 1365 **Metacognitive Reasoning System Learning (MeRSL)**, a RL method designed to further enhance  
 1366 emergent metacognition system collaboration.  
 1367

1368 **Enhancing Metacognition Loop for Emergent Metacognition System.** The Prompt-Driven  
 1369 Metacognitive Reasoning System (§ 5.1) demonstrates that explicit role-playing effectively elicits  
 1370 latent reasoning capabilities. However, this “emergent” success relies heavily on the inherent  
 1371 instruction-following and in-context learning abilities of large-scale foundation models (e.g.,  
 1372 Gemini-2.5-Pro). Smaller models, despite possessing the foundational metacognitive signals  
 1373 identified in § 5.2, often struggle to maintain such complex functional loops via prompting alone.  
 1374 They lack the stability to coordinate distinct roles purely through context.

1375 This limitation motivates us to ask: *Can we optimize the collaborative dynamics between metacog-  
 1376 nitive and cognitive roles via training?* Building on recent findings that RL can enhance sub-  
 1377 components of metacognition system (Yang & Thomason, 2025; Wan et al., 2025), MeRSL aims  
 1378 to bridge this gap. By formalizing the interaction between roles as a trainable cooperative game, we  
 1379 can improve our metacognition system.

1380 MeRSL structures the reasoning process as a hierarchical interaction between two distinct policy  
 1381 levels. We decompose the optimization into two sequential stages corresponding to the cognitive  
 1382 loop in Fig. 7: (1) **Stage 1**: Optimizing the collaboration between a Meta-Agent (assessing difficulty  
 1383 and generating plans) and a Reasoning-Agent (executing the solution). (2) **Stage 2**: Optimizing the  
 1384 interaction between a Monitor-Agent (evaluating confidence) and a Strategy-Agent (adjusting the  
 1385 path).

1386 **Stage 1.** We decouple the generation process into a high-level meta-policy  $\pi_h$  and a low-level  
 1387 reasoning-policy  $\pi_l$ . Here, we leverage SFT as a cold start to introduce the high-level agent to  
 1388 various difficulty formats it can utilize.

1389 Formally, The high-level policy ( $\pi_h$ ) first conditions on the input  $\mathbf{x}$  to generate a metacognitive  
 1390 directive  $\mathbf{m}$  (e.g., <difficulty>...<plan>...). Subsequently, the low-level policy  $\pi_l$  con-  
 1391 ditions on both the input and this directive to produce the reasoning trajectory  $\mathbf{y}$ . The generation  
 1392 process is formulated as:

$$\mathbf{y} \sim \pi_l(\mathbf{y} \mid \mathbf{x}, \mathbf{m}) \pi_h(\mathbf{m} \mid \mathbf{x}). \quad (8)$$

1393  
 1394  
 1395  
 1396 *Optimization.* During training, suppose  $\theta_h$  and  $\theta_l$  denote the parameters for the high-level and low-  
 1397 level agents, respectively. The joint system policy  $\pi_{(\theta_h, \theta_l)}$  is formulated as:

$$\mathbf{y} \sim \pi_{(\theta_h, \theta_l)}(\mathbf{y} \mid \mathbf{x}) := \pi_{\theta_l}(\mathbf{y} \mid \mathbf{x}, \mathbf{m}) \cdot \pi_{\theta_h}(\mathbf{m} \mid \mathbf{x}), \quad (9)$$

1400 The objective is to maximize the expected reward  $R(\mathbf{y}, \mathbf{y}^*)$ , which is defined as:  
 1401

$$\mathcal{J}(\theta_h, \theta_l) = \mathbb{E}_{\mathbf{x}, \mathbf{y}^*} \mathbb{E}_{\mathbf{y} \sim \pi_{(\theta_h, \theta_l)}} R(\mathbf{y}, \mathbf{y}^*). \quad (10)$$

1404 We adopt an iterative optimization strategy where each agent maximizes their respective rewards  
 1405 independently. The optimization is decoupled as follows:

$$\theta_h^* = \arg \max_{\theta_h} \mathbb{E}_{(\mathbf{x}, \mathbf{y}^*) \sim \mathcal{D}, \mathbf{m} \sim \pi_{\theta_h}, \mathbf{y} \sim \pi_{\theta_h^*}} [R_h(\mathbf{m}, \mathbf{y}, \mathbf{y}^*)], \quad (11)$$

$$\theta_l^*(\theta_h) = \arg \max_{\theta_l} \mathbb{E}_{(\mathbf{x}, \mathbf{y}^*) \sim \mathcal{D}, \mathbf{m} \sim \pi_{\theta_h}, \mathbf{y} \sim \pi_{\theta_l}} [R_l(\mathbf{m}, \mathbf{y}, \mathbf{y}^*)], \quad (12)$$

1410 where  $R_h$  and  $R_l$  are policies' individual reward functions.

1411 *Reward design.* Concretely, the high-level agent is trained to (i) produce metacognitive  
 1412 directives that lead to consistent low-level solutions, (ii) obey the required  
 1413 “<difficulty>...<plan>...” format without leaking final answers, and (iii) cor-  
 1414 rectly classify the pre-bucketed difficulty. The low-level agent is trained to (i) solve the problem  
 1415 correctly under the given directive, (ii) follow the required reasoning format, and (iii) allocate  
 1416 reasoning budget conditioned on the predicted difficulty.

1417 For high-level agent, the meta-policy  $\pi_h$  is evaluated based on the stability of the downstream exe-  
 1418 cution, the accuracy of its self-assessment, and structural compliance.

1420 *Consistency Reward ( $R_{\text{cons}}$ ).* Following the intuition that a high-quality plan should lead to unam-  
 1421 biguous execution, we reward the meta-agent for reducing the entropy of the low-level agent’s output  
 1422 distribution. Given a directive  $\mathbf{m}$ , we sample  $K$  rollouts  $\{\mathbf{y}^{(k)}\}_{k=1}^K$  from  $\pi_l(\cdot | \mathbf{x}, \mathbf{m})$  and extract  
 1423 their final answers  $a^{(k)}$ . The consistency reward is defined as the empirical majority vote ratio:

$$R_{\text{cons}}(\mathbf{m}) = \max_a \frac{1}{K} \sum_{k=1}^K \mathbb{I}[a^{(k)} = a]. \quad (13)$$

1428 This objective encourages  $\pi_h$  to generate directives that guide  $\pi_l$  toward a stable solution mode,  
 1429 mitigating reasoning variance.

1430 *Difficulty Calibration Reward ( $R_{\text{diff}}$ ).* To ground the agent’s metacognition in objective standards, we  
 1431 introduce a calibration term. Let  $d^* \in \{\text{easy}, \text{medium}, \text{hard}\}$  denote the ground-truth difficulty  
 1432 bucket, and  $\hat{d} = g(\mathbf{m})$  be the predicted difficulty parsed from the directive. We apply a binary  
 1433 reward for correct classification:

$$R_{\text{diff}}(\mathbf{m}) = \mathbb{I}[\hat{d} = d^*]. \quad (14)$$

1436 Weighted by  $\lambda = 0.2$ , this term explicitly aligns the model’s internal assessment with the dataset’s  
 1437 complexity distribution without dominating the planning objective.

1439 *Format Regularization ( $R_{\text{fmt}}^h$ ).* We enforce structural adherence via a format reward  $R_{\text{fmt}}^h$ . The agent  
 1440 receives a positive signal for correctly generating the <difficulty> and <plan> tags. Cru-  
 1441 cially, to enforce the abstraction boundary between planning and solving, the agent incurs a severe  
 1442 penalty if it generates solution-specific artifacts (e.g., \boxed{...}) within the planning phase.

1443 The total high-level reward is aggregated as:

$$R_h(\mathbf{m}, \mathbf{y}, \mathbf{y}^*) = R_{\text{cons}}(\mathbf{m}) + \lambda R_{\text{diff}}(\mathbf{m}) + R_{\text{fmt}}^h(\mathbf{m}). \quad (15)$$

1447 For low-level agent, the reasoning policy  $\pi_l$  focuses on solving the problem correctly, but with a  
 1448 novel constraint: it must allocate computational resources commensurate with the assessed diffi-  
 1449 culty.

1451 *Correctness Reward ( $R_{\text{cor}}$ ).* The primary learning signal remains the binary correctness of the final  
 1452 answer against the ground truth  $f(\mathbf{y}^*)$ :

$$R_{\text{cor}}(\mathbf{y}, \mathbf{y}^*) = \mathbb{I}[f(\mathbf{y}) = f(\mathbf{y}^*)]. \quad (16)$$

1456 *Difficulty-Adaptive Length Penalty ( $R_{\text{len}}$ ).* To realize metacognitive control over the reasoning bud-  
 1457 get, we introduce a difficulty-conditioned length penalty. Let  $T(\mathbf{y})$  be the token length of the rea-  
 1458 soning trace and  $T_{\text{ref}}$  be a reference length (e.g., the average CoT length in the training set). The

1458 penalty is modulated by the predicted difficulty  $\hat{d}$ :  
 1459

$$1460 \quad R_{\text{len}}(\mathbf{m}, \mathbf{y}) = -\alpha(\hat{d}) \cdot \frac{T(\mathbf{y})}{T_{\text{ref}}}, \quad \text{where} \quad \alpha(\hat{d}) = \begin{cases} 0.2, & \text{if } \hat{d} = \text{easy}, \\ 0.1, & \text{if } \hat{d} = \text{medium}, \\ 0, & \text{if } \hat{d} = \text{hard}. \end{cases} \quad (17)$$

1461  
 1462  
 1463

1464 This piecewise function  $\alpha(\hat{d})$  serves as a soft constraint: it discourages verbosity for simple prob-  
 1465 lems (efficiency) while effectively uncapping the reasoning budget for hard problems (exploration),  
 1466 aligning compute allocation with the task’s intrinsic complexity.  
 1467

1468 *Format Regularization* ( $R_{\text{fmt}}^l$ ). A lightweight term  $R_{\text{fmt}}^l$  rewards the presence of valid reasoning  
 1469 delimiters (e.g., `<think>...</think>`) and penalizes malformed traces, ensuring the stability  
 1470 of the parsing and evaluation pipeline.

1471 The total low-level reward is summarized as:  
 1472

$$R_l(\mathbf{m}, \mathbf{y}, \mathbf{y}^*) = R_{\text{cor}}(\mathbf{y}, \mathbf{y}^*) + R_{\text{len}}(\mathbf{m}, \mathbf{y}) + R_{\text{fmt}}^l(\mathbf{y}). \quad (18)$$

1473  
 1474

1475 **Stage 2.** After Stage 1 establishes the meta–reasoning hierarchy, we further optimize the collabora-  
 1476 tion between the *Monitor-Agent* (confidence/error evaluation) and the *Strategy-Agent* (trajectory  
 1477 adjustment). Concretely, we instantiate three role-conditioned policies from the same underlying  
 1478 LRM: a reasoning policy  $\pi_{\theta_l}$ , a monitor policy  $\pi_{\theta_m}$ , and a strategy policy  $\pi_{\theta_s}$ . Given the input  $\mathbf{x}$   
 1479 and the metacognitive directive  $\mathbf{m}$  produced in Stage 1, the single-round correction process is:  
 1480

$$1481 \quad \mathbf{y} \sim \pi_{\theta_l}(\mathbf{y} \mid \mathbf{x}, \mathbf{m}), \quad \mathbf{e} \sim \pi_{\theta_m}(\mathbf{e} \mid \mathbf{x}, \mathbf{m}, \mathbf{y}), \quad \mathbf{y}' \sim \pi_{\theta_s}(\mathbf{y}' \mid \mathbf{x}, \mathbf{m}, \mathbf{y}, \mathbf{e}), \quad (19)$$

1482 where  $\mathbf{y}$  is the initial reasoning trajectory,  $\mathbf{e}$  is the monitor report, and  $\mathbf{y}'$  is the corrected trajectory.  
 1483 We perform only *one* Monitor→Strategy correction round in training.

1484 *Grouped rollouts.* For each training instance  $(\mathbf{x}, \mathbf{m}, \mathbf{y}^*)$ , we first sample  $G$  rollout:  
 1485

$$1486 \quad \{\mathbf{y}^{(i)}\}_{i=1}^G, \quad \mathbf{y}^{(i)} \sim \pi_{\theta_l}(\cdot \mid \mathbf{x}, \mathbf{m}). \quad (20)$$

1487

1488 Each  $\mathbf{y}^{(i)}$  is evaluated by a correctness reward defined in Eq. 16

1489 *Priority-Guided Seed Sampling.* To train the Monitor-Agent on informative cases, we select  $K$  seeds  
 1490 from the reasoning group using *Priority-Guided Seed Sampling* (PGSS). PGSS prioritizes incorrect  
 1491 trajectories: (1) if at least  $K$  trajectories have  $R_l^{(i)} = 0$ , we uniformly sample  $K$  seeds from them;  
 1492 (2) otherwise, we take all incorrect trajectories and uniformly sample the remaining seeds from the  
 1493 correct ones. Denote the selected reasoning seeds as  $\{\tilde{\mathbf{y}}^{(k)}\}_{k=1}^K$ .

1494 For each seed  $\tilde{\mathbf{y}}^{(k)}$ , the Monitor-Agent samples a report group of size  $G$ :  
 1495

$$1496 \quad \{\mathbf{e}^{(k,j)}\}_{j=1}^G, \quad \mathbf{e}^{(k,j)} \sim \pi_{\theta_m}(\cdot \mid \mathbf{x}, \mathbf{m}, \tilde{\mathbf{y}}^{(k)}). \quad (21)$$

1497

1498 Each report contains a judgment  $\hat{z}^{(k,j)} \in \{\text{OK}, \text{ERROR}\}$  indicating whether the seed trajectory is  
 1499 believed to be correct.

1500 *Monitor reward.* Let  $z^{(k)} = \mathbb{I}[f(\tilde{\mathbf{y}}^{(k)}) = f(\mathbf{y}^*)]$  be the true correctness of the seed. We define a  
 1501 local monitoring reward,  
 1502

$$1503 \quad R_{\text{loc}}^{(k,j)} = \begin{cases} +1, & \hat{z}^{(k,j)} = \text{OK} \text{ and } z^{(k)} = 1, \\ +1, & \hat{z}^{(k,j)} = \text{ERROR} \text{ and } z^{(k)} = 0, \\ -1, & \text{otherwise,} \end{cases} \quad (22)$$

1504  
 1505

1506 encouraging accurate confidence/error judgments.  
 1507

1508 We further introduce a *correction-gain reward* to align monitoring with downstream usefulness.  
 1509 From the reports that predict `ERROR`, we select up to  $K$  report seeds using *Reliability-Guided*  
 1510 *Report Sampling* (RGRS): we first prioritize reports with  $R_{\text{loc}}^{(k,j)} = +1$ , and fill any remaining slots  
 1511 with other `ERROR` reports. If no `ERROR` report is produced, the correction branch terminates early.  
 Denote a selected report seed as  $\tilde{\mathbf{e}}^{(k,m)}$ .

1512 Table 30: Stage-wise results with a single correction round at inference.  
1513  
1514

Dataset	Prompt-based Loop	Stage 1	Stage 2	Stage 1 + Stage 2
MATH500	75.3	76.2	77.2	<b>79.4</b>
GSM8K	92.2	91.8	92.6	92.6

1515 For each selected  $\tilde{\mathbf{e}}^{(k,m)}$ , the Strategy-Agent samples  $G$  corrected trajectories:  
1516  
1517

$$\{\mathbf{y}'^{(k,m,\ell)}\}_{\ell=1}^G, \quad \mathbf{y}'^{(k,m,\ell)} \sim \pi_{\theta_s}(\cdot \mid \mathbf{x}, \mathbf{m}, \tilde{\mathbf{y}}^{(k)}, \tilde{\mathbf{e}}^{(k,m)}), \quad (23)$$

1518 with correctness reward  
1519  
1520

$$R_s^{(k,m,\ell)} = \mathbb{I}[f(\mathbf{y}'^{(k,m,\ell)}) = f(\mathbf{y}^*)]. \quad (24)$$

1521 Let the correction success ratio be  
1522  
1523

$$p^{(k,m)} = \frac{1}{G} \sum_{\ell=1}^G R_s^{(k,m,\ell)}. \quad (25)$$

1524 The gain reward for the selected monitor report is  
1525  
1526

$$R_{\text{gain}}^{(k,m)} = \begin{cases} p^{(k,m)}, & z^{(k)} = 0 \text{ (incorrect seed improved)}, \\ -(1 - p^{(k,m)}), & z^{(k)} = 1 \text{ (correct seed degraded)}. \end{cases} \quad (26)$$

1527 Unselected reports receive  $R_{\text{gain}} = 0$ . The total monitor reward is  
1528  
1529

$$R_m^{(k,j)} = R_{\text{loc}}^{(k,j)} + R_{\text{gain}}^{(k,j)}. \quad (27)$$

1530 *Optimization.* We update the three role policies using GRPO. For each group of size  $G$  generated  
1531 under the same conditioning context, we compute normalized intra-group advantages  
1532

$$\hat{A}_i = \frac{R_i - \text{mean}(R_{1:G})}{\text{std}(R_{1:G}) + \epsilon}, \quad (28)$$

1533 and apply a clipped policy-gradient update with KL regularization to the corresponding role. Reasoning  
1534 groups use rewards  $\{R_l^{(i)}\}$ , monitor groups use  $\{R_m^{(k,j)}\}$ , and strategy groups use  $\{R_s^{(k,m,\ell)}\}$ .  
1535 This joint training encourages the Monitor-Agent to produce diagnostically useful feedback and the  
1536 Strategy-Agent to reliably correct low-confidence trajectories, thereby strengthening the emergent  
1537 metacognitive loop under limited model capacity.  
15381539 **Details.** We conduct training on `Qwen2.5-7B-Instruct` using the MATH dataset. We optimize  
1540 the model with Adam Optimizer using a constant learning rate of  $1e-6$ . During rollout, the prompt  
1541 batch size is set to 128, and we sample  $G = 8$  responses for each prompt. The sampling temperature  
1542 is 1 with  $\text{top-}p = 1.0$  and  $\text{top-}k = -1$ . We cap the maximum response length at 2048 tokens.  
1543 For stage 1, we set  $\lambda = 0.2$  and  $T_{\text{ref}} = 1024$ . For stage 2, we set seeds  $K = 2$ .  
15441545 **Results.** Under the constrained inference setting where only a single correction round is allowed, we  
1546 evaluate the effectiveness of Stage 1, Stage 2, and their combination. All experiments are conducted  
1547 with `Qwen2.5-7B-Instruct` on two standard mathematical reasoning benchmarks: MATH500  
1548 and GSM8K. We use the instruct model to ensure better instruction-following behavior.  
15491550 On GSM8K, Stage 1 yields a mild degradation. We conjecture this is because GSM8K problems are  
1551 relatively simple and short-horizon; explicit planning and difficulty self-assessment provide limited  
1552 additional benefit, while introducing a small overhead that can occasionally distract the low-level  
1553 execution. In contrast, Stage 2 consistently improves GSM8K, suggesting that even for easier problems,  
1554 a monitor-strategy correction step can fix local arithmetic slips. Combining Stage 1 and Stage  
1555 2 maintains the gain from Stage 2 without further improvement, implying that Stage 2 dominates  
1556 the achievable headroom under the one-round inference constraint. Overall, these results demon-  
1557 strate that our trainable metacognitive collaboration yields robust gains on difficult mathematical  
1558 reasoning, while remaining effective under strict inference budgets.  
1559

1566

Table 31: Metacognitive Reasoning System: Prompt of Step 1.

1567

1568

1569

**System Prompt:**

1570 You are an elite mathematical strategist and analyst. Your  
 1571 primary function is to perform a deep Metacognitive Analysis of  
 1572 complex mathematical problems. You are to deconstruct the problem  
 1573 into its core components, identify underlying principles, and then  
 1574 formulate a high-level, executable strategic plan.

1575 Your task is to produce a Metacognitive Analysis of the following  
 1576 problem. You must NOT provide a final solution or perform  
 1577 detailed calculations.

1578 **### Core Principles**

1579 \* **\*\*Analytical Depth:\*\*** Your analysis must go beyond a  
 1580 surface-level reading. Identify the mathematical field, key  
 1581 concepts, constraints, and the explicit goal.

1582 \* **\*\*Strategic Foresight:\*\*** Your plan should be a viable path to  
 1583 a solution. This includes anticipating potential difficulties,  
 1584 identifying necessary lemmas, and choosing the most promising  
 1585 approach.

1586 \* **\*\*Clarity and Brevity:\*\*** The analysis and plan must be clear,  
 1587 concise, and easily understood by another mathematical expert who  
 1588 will execute it.

1589 **User Prompt:**

1590 **### Your Task**

1591 **\*\*Problem:\*\***  
 =====

1592 {problem\_statement}  
 =====

1593 **\*\*Metacognitive Analysis:\*\***

1594 **\*\*1. Problem Deconstruction:\*\***

1595 \* **\*\*Mathematical Domain:\*\*** Identify the primary field(s) of  
 1596 mathematics involved (e.g., Number Theory, Combinatorics,  
 1597 Euclidean Geometry).

1598 \* **\*\*Given Conditions & Constraints:\*\*** List all the premises,  
 1599 conditions, and constraints provided in the problem statement in  
 1600 a structured format.

1601 \* **\*\*Objective:\*\*** State the precise question to be answered or the  
 1602 proposition to be proven.

1603 **\*\*2. Strategic Solution Plan (Method Sketch):\*\***

1604 Present a high-level, conceptual outline of your proposed solution  
 1605 path. This sketch should enable an expert to grasp the entire  
 1606 logical flow of the argument without needing the full details. It  
 1607 must include:

1608 \* **\*\*Overall Strategy Narrative:\*\*** A brief description of the core  
 1609 idea behind your approach (e.g., "We will use proof by induction,"  
 1610 "The strategy is to establish a coordinate system and use analytic  
 1611 geometry," "We will prove the contrapositive by assuming...").

1612 \* **\*\*Key Lemmas and Intermediate Results:\*\*** State the full and  
 1613 precise mathematical formulations of any key lemmas or theorems  
 1614 you plan to prove or apply. These are the major milestones of the  
 1615 proof.

1616 \* **\*\*Logical Skeleton:\*\*** If applicable, describe the key  
 1617 constructions, case splits, or transformations that form the  
 1618 backbone of your argument.

1619 \* **\*\*Potential Challenges & Pitfalls:\*\*** Briefly note any steps  
 1620 that might be particularly tricky, prone to error, or require a  
 1621 non-obvious insight.

1622 **### Negative Constraints**

1623 \* **\*\*DO NOT\*\*** write the full, step-by-step solution.

1624 \* **\*\*DO NOT\*\*** perform detailed algebraic manipulations or numerical  
 1625 calculations.

1626 \* Your output should be strictly limited to the analysis and  
 1627 strategic plan as outlined above.

1620  
1621  
1622  
1623  
1624  
1625  
1626

Table 32: Metacognitive Reasoning System: Prompt of Step 2.

1627 **System Prompt:**  
 1628 You are an exceptionally rigorous mathematical solver. Your  
 1629 sole purpose is to take a pre-defined strategic plan and execute  
 1630 it with absolute precision and logical soundness. You must not  
 1631 deviate from, question, or reinterpret the provided plan.  
 1632 Your task is to produce a complete and formally justified solution  
 1633 to the following mathematical problem, strictly following the  
 1634 'Solution Plan'.  
 1635 **### Core Principles**  
 1636 \* **\*\*Rigor is Paramount:\*\*** Your primary goal is to produce a  
 1637 complete and rigorously justified solution. Every step in your  
 1638 solution must be logically sound and clearly explained. A correct  
 1639 final answer derived from flawed or incomplete reasoning is  
 1640 considered a failure.  
 1641 \* **\*\*Unyielding Adherence to Plan:\*\*** You MUST strictly follow  
 1642 the logical flow, lemmas, and constructions laid out in the  
 1643 'Solution Plan'. Do not introduce new methods, skip steps, or  
 1644 alter the proposed strategy in any way. Your role is execution,  
 1645 not creation.  
 1646 \* **\*\*Honesty About Completeness:\*\*** If you cannot find a complete  
 1647 solution following the plan, you must **\*\*not\*\*** guess or create a  
 1648 solution that appears correct but contains hidden flaws. Instead,  
 1649 you should present only the significant partial results that you  
 1650 can rigorously prove by following the plan.  
 1651 **User Prompt:**  
 1652 **### Your Task**  
 1653 **\*\*Problem:\*\***  
 1654 =====  
 1655 {problem\_statement}  
 1656 =====  
 1657 **\*\*Solution Plan:\*\***  
 1658 =====  
 1659 {step1.output}  
 1660 =====  
 1661 **\*\*Detailed Solution:\*\***  
 1662 Present the full, step-by-step mathematical proof, meticulously  
 1663 following the guidance of the 'Solution Plan'. Each step must be  
 1664 logically justified and clearly explained. The level of detail  
 1665 should be sufficient for an expert to verify the correctness of  
 1666 your reasoning without needing to fill in any gaps. This section  
 1667 must contain ONLY the complete, rigorous proof, free of any  
 1668 internal commentary, alternative approaches, or failed attempts.  
 1669 \* **\*\*Use TeX for All Mathematics:\*\*** All mathematical variables,  
 1670 expressions, and relations must be enclosed in TeX delimiters  
 1671 (e.g., 'Let  $n$  be an integer.').  
 1672 Your step-by-step reasoning, strictly following the plan, begins  
 1673 here...  
 1674 **### Final Answer**  
 1675 After completing the detailed solution, state the final answer  
 1676 within `\boxed{}`.  
 1677  
 1678  
 1679  
 1680  
 1681  
 1682  
 1683  
 1684  
 1685  
 1686  
 1687  
 1688  
 1689  
 1690  
 1691  
 1692  
 1693  
 1694  
 1695  
 1696  
 1697  
 1698  
 1699  
 1700  
 1701  
 1702  
 1703

1674

1675

1676

1677

Table 33: Metacognitive Reasoning System: Prompt of Step 3 (Part 1).

1678

1679

**System Prompt:**

You are an expert mathematician and a meticulous grader for AIME-level computational problems. Your primary task is to rigorously verify the provided solution's \*\*computational reasoning and numeric correctness\*\*. A solution is to be judged correct \*\*only if every step that affects the numeric outcome is correct and sufficiently justified.\*\* A solution that reaches a correct final integer answer via arithmetic slips, incorrect algebraic manipulations, unverified casework, counting mistakes, or hidden assumptions must be flagged as incorrect or incomplete.

### Instructions ###

\*\*1. Core Instructions\*\*

- \* Your sole task is to identify and report all issues in the provided solution. You must act strictly as a \*\*verifier\*\*, NOT a solver.
- \* You must \*\*NOT attempt to correct, fix, or complete\*\* any errors or missing arguments.
- \* Perform a \*\*step-by-step\*\* check of the entire solution and produce a \*\*Detailed Verification Log\*\*. For each step:
  - \* If the step is correct, state briefly that it is correct.
  - \* If the step contains an issue, explain the error and classify it (see section 2).

\*\*2. How to Handle Issues in the Solution\*\*

All issues must be classified into one of the following categories:

- \* \*\*a. Critical Error:\*\*
- \* Definition: Any error that changes or potentially invalidates the numeric result. Examples include arithmetic mistakes, wrong algebraic transformations, misapplied formulas, incorrect combinatorial counts, invalid casework, or unjustified approximations that affect the integer outcome.
- \* \*\*Procedure:\*\*
- \* Point out the exact error and explain why it invalidates the reasoning.
- \* Do \*\*not\*\* check further steps that rely on this error.
- \* You may still check other independent parts of the solution.
- \* \*\*b. Justification Gap:\*\* Definition: Steps where the stated conclusion might be correct, but the reasoning is incomplete or not justified at AIME level.
- \* \*\*Procedure:\*\*
- \* Point out the missing justification.
- \* Explicitly state that you will assume the step's conclusion holds for the sake of checking subsequent steps.

\*\*3. Output Format\*\*

Your response MUST be structured into two main sections: a \*\*Summary\*\* followed by the \*\*Detailed Verification Log\*\*.

- \* \*\*a. Summary\*\* \*\*Final Verdict:\*\* One clear sentence declaring overall validity (e.g., "The solution is correct," "The solution contains a Critical Error and is therefore invalid," or "The solution contains several Justification Gaps.").
- \* \*\*List of Findings:\*\* A bulleted list of every issue found. For each finding include:
  - \* \*\*Location:\*\* A direct quote of the key phrase or equation.
  - \* \*\*Issue:\*\* Short description and classification (\*\*Critical Error\*\* or \*\*Justification Gap\*\*).

1725

1726

1727

Table 34: Metacognitive Reasoning System: Prompt of Step 3 (Part 2).

```

1728
1729
1730 * **b. Detailed Verification Log*** Provide a step-by-step
1731 verification.
1732 * Quote the relevant part of the solution before your check.
1733 * State clearly: **Correct**, **Critical Error**, or
1734 **Justification Gap**.
1735 * Do **not** supply corrections or alternative methods | only
1736 report the issues.
1737 * **Important:**
1738 - Do not propose fixes or alternative solutions.
1739 - Do not attempt to supply missing reasoning.
1740 - Only check and report correctness of what is written.

1741 User Prompt:
1742 ##### Your Task
1743 ***Original Problem:***
1744 =====
1745 {problem_statement}
1746 =====
1747 ***Current Solution:***
1748 =====
1749 {last_solution}
1750 =====

1751 ##### Monitoring Task Reminder #####
1752 Your task is to act as an math grader. Now, generate the
1753 ***summary*** and the **step-by-step verification log** for the
1754 solution above. In your log, justify each correct step and
1755 explain in detail any errors or justification gaps you find, as
1756 specified in the instructions above."
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781

```

### A.7 § 5.2. METACOGNITIVE REASONING MODELS: MORE DETAILS

**Implementation Details.** For the online model APIs, we utilized their respective official endpoints with default temperature settings. The maximum generation length was also kept at its default value, and no other parameters were modified.

For the training of open-source models, all experiments were conducted on a server equipped with 8×NVIDIA H800 GPU (80GB). The maximum generation length was set to 8192 for the GSM8K and MATH datasets, and 16384 for all other datasets. The temperature was set to 1.0, and the random seed was fixed to 42 for reproducibility.

For SFT training, we adopt full-parameter fine-tuning with learning rate  $r = 1 \times 10^{-5}$  and  $batchsize = 32$ . To construct training data, we first define the difficulty level for each sample from the train set of GSM8K and MATH: Easy (GSM8K and MATH lv.1), Medium (MATH lv.1-4), and Hard (MATH lv.3-5). Second, to obtain high-level task decomposition, we utilize Gemini-2.5-Pro with the prompt in Appendix A.4. The final pattern of training samples in cold-start is in the form of: ``<difficulty> level </difficulty> <plan> decomposition </plan> <think> CoT </think> answer''.

For RL training, we write our code based on the open-source Verl framework. Training settings are listed in Tab. 40. During inference, we used a standard prompt without any task-specific engineering during inference for all models (Base, GRPO, and Ours): “*Please reason step by step and provide your final answer within \boxed{}*.”

**Discussion on Signal Selection in RL** In our MRM RL training, we conducted an ablation study on signal selection. We compared our default *Token Confidence* against (1) *Sentence Confidence*, the average confidence of the current sentence. *Sentence Entropy*, The mean entropy of the sentence token. *Random Forking*, A control baseline where forking occurs at random positions until the target group size is reached, decoupling exploration from uncertainty.

1782  
1783  
1784  
1785  
1786  
1787

Table 35: Metacognitive Reasoning System: Prompt of Step 4 (Part 1).

1788  
1789  
1790  
1791  
1792  
1793  
1794  
1795  
1796  
1797  
1798  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1830  
1831  
1832  
1833  
1834  
1835

**System Prompt:**

You are an expert mathematician and a careful corrector for AIME-level computational problems.  
You will be given three inputs:  
1) The Original Problem,  
2) The current Solution,  
3) A Verification Log (from a previous check), which labels each step as Correct / Justification Gap / Critical Error, and provides short notes.  
### Your Task ###  
Using the Verification Log, \*\*step by step correct the Original Solution\*\*.  
- If a step is labeled \*\*Correct\*\*, keep it unchanged (you may lightly reformat for clarity).  
- If a step is labeled \*\*Justification Gap\*\*, supply the missing justification or intermediate calculations, enough for AIME-level rigor.  
- If a step is labeled \*\*Critical Error\*\*, replace it with a correct mathematical step (with explicit computations or reasoning) and update all dependent later steps accordingly.  
- Do \*\*not\*\* introduce new solution paths, alternative methods, or multiple approaches. Only repair the given solution chain.  
### Output Format ###  
1. \*\*Correction Summary\*\*  
- A single sentence declaring whether the solution has been fully corrected and what the final answer is.  
- Example: \The solution has been fully corrected. Final Answer = 70."  
- Or, if not possible: \The solution cannot be fully corrected due to missing information in step X."  
2. \*\*Correction Log\*\*  
For each relevant step (especially those flagged in the Verification Log), provide an entry with:  
- \*\*Quoted Step:\*\* The original line/equation (quoted or in a code block).  
- \*\*Verification Label:\*\* Correct / Justification Gap / Critical Error.  
- \*\*Correction / Action:\*\*  
\* If Correct → \Unchanged | correct."  
\* If Justification Gap → Provide the missing computation/derivation briefly, ending with \Filled gap."  
\* If Critical Error → Provide the corrected computation/derivation, briefly note why the original was wrong, and end with \Corrected."  
- If a step's correction affects later steps, explicitly note \Affects subsequent steps: Yes/No."  
3. \*\*Full Corrected Solution\*\*  
- Present the entire solution in a clean, continuous write-up, combining unchanged and corrected steps.  
- Show all necessary algebra, arithmetic, or combinatorial reasoning clearly.  
- After completing the detailed solution, state the final answer within \boxed{}.

Table 36: Metacognitive Reasoning System: Prompt of Step 4 (Part 2).

```

1836
1837
1838 User Prompt:
1839     ### Your Task
1840     ***Original Problem:**
1841     =====
1842     {problem_statement}
1843     =====
1844     **Current Solution:**
1845     =====
1846     {last_solution}
1847     =====
1848     **Verification Log:**
1849     =====
1850     {monitor_output}
1851     =====

```

Table 37: Ablation study on RL performance with different monitoring signals. **Random** denotes random forking strategies.

Metric Type	MATH500 (Pass@1)	GSM8K (Pass@1)
Qwen2.5-Math-7B (Base)	64.0%	70.3%
<i>Baseline (GRPO)</i>	71.6%	75.9%
Random Forking	66.3%	71.4%
Sentence Confidence	72.4%	79.1%
Sentence Entropy	73.0%	78.8%
<b>Token Confidence (Ours)</b>	<b>80.2%</b>	<b>85.5%</b>

**Results.** As shown in Tab. 37, both sentence-level confidence and entropy outperform the standard GRPO baseline, indicating that the uncertainty estimation is robust across different signal definitions. We also find the random setting leads to performance degradation compared to GRPO. This negative result is crucial: it confirms that simply increasing exploration diversity is insufficient and can be detrimental. Overall, these results validate that the primary driver of improvement is our metacognitive control mechanism.

Table 38: Performance (%) on Out-of-Distribution (OOD) Benchmarks. **Ours (SFT+RL)** demonstrates superior generalization capabilities compared to baselines.

Method	ARC-c	GPQA-Diamond	MMLU-Pro	LiveCodeBench
Base Model (Qwen2.5-Math-7B)	18.2	33.8	37.4	28.9
GRPO (Standard RL)	29.8	39.2	42.1	35.5
<b>Ours (SFT+RL)</b>	<b>33.2</b>	<b>46.5</b>	<b>47.6</b>	<b>40.4</b>

**Out-of-Distribution (OOD) Generalization.** To validate generalization, we evaluated our fine-tuned model (*SFT+RL*) against the *Qwen2.5-Math-7B* and the standard RL baseline (*GRPO*) on four diverse OOD benchmarks: **ARC-c** (Clark et al., 2018): A challenging open-domain reasoning benchmark. **GPQA-DIAMOND**: A graduate-level science benchmark covering biology, physics, and chemistry. **MMLU-PRO** (Wang et al., 2024): A comprehensive benchmark focusing on complex reasoning across diverse academic subjects. **LIVECODEBENCH** (Jain et al., 2024): A holistic evaluation for code generation, representing a significant domain shift from mathematics. To avoid contamination, we shuffled the multiple-choice options for all QA tasks.

**Results.** The performance comparison is presented in Tab. 38. Our method consistently outperforms both the Base Model and the GRPO baseline across all benchmarks. Specifically, our method achieves substantial gains on scientific reasoning and general academic reasoning. This confirms that the internalized metacognitive skills like *planning* are not merely overfitting to math problems but can transfer effectively to general reasoning tasks. On **LIVECODEBENCH**, our method shows

1890 a modest but meaningful improvement. We attribute the relative difficulty of this task to the base  
 1891 model architecture (Qwen2.5-Math), which is specialized for mathematics rather than coding.  
 1892

1893 **Integration with Self-Improvement Paradigms.** A critical extension regarding our framework is  
 1894 its relationship with self-training methods like STaR (Zelikman et al., 2024). we conducted a con-  
 1895 trolled experiment combining our Metacognitive Schema with the STaR iterative loop. We utilized  
 1896 Qwen2.5-Math-7B as the base model and evaluated performance on the MATH-500 benchmark  
 1897 across multiple iterations. The comparison settings were *Standard STaR*: Iteratively fine-tuning on  
 1898 standard rationales:  $Q \rightarrow \text{CoT} \rightarrow A$ , and our *Meta-STaR*: Iteratively fine-tuning on structured  
 1899 metacognitive traces:  $Q \rightarrow \langle \text{difficulty} \rangle \rightarrow \langle \text{plan} \rangle \rightarrow \langle \text{think} \rangle \rightarrow A$ .  
 1900

1901 Table 39: Performance comparison on MATH-500 across bootstrapping iterations. **Meta-STaR**  
 1902 demonstrates a faster rate of improvement compared to standard STaR.  
 1903

Method	Iteration 0 (Base)	Iteration 2	Iteration 4
Qwen2.5-Math-7B	64.0%	-	-
Standard STaR	64.0%	65.7%	67.0%
<b>Meta-STaR (Ours)</b>	64.0%	<b>68.6%</b>	<b>71.2%</b>

1904 **Results.** As detailed in Tab. 39, while standard STaR yields consistent improvements, applying the  
 1905 STaR algorithm to our metacognitive schema results in significantly higher gains. This substantial  
 1906 gap suggests that incorporating metacognitive signals makes the bootstrapping process more effi-  
 1907 cient. By enforcing explicit planning and self-assessment, the model generates higher-quality ratio-  
 1908 nales during the exploration phase, thereby creating superior training data for subsequent iterations.  
 1909 This confirms that our structured metacognitive paradigm can effectively serve as a foundational  
 1910 architecture for advanced self-improvement algorithms.  
 1911

## 1912 A.8 STATEMENT ON THE USE OF LARGE LANGUAGE MODELS (LLMs)

1913 **LLMs as the Subject of Research.** One of the core components of our research involves the  
 1914 investigation and evaluation of the reasoning capabilities of current open-source and closed-source  
 1915 Large Language Models (LLMs). As such, a number of LLMs are explicitly named and analyzed  
 1916 within this paper (as detailed in § 4). In this capacity, they serve as the objects of our study.  
 1917

1918 **LLMs as an Assistive Tool.** In the preparation of this manuscript, the use of LLMs is limited to  
 1919 polishing the text for grammatical correctness, spelling, and clarity of expression. The LLMs were  
 1920 not used to generate any core research ideas, experimental designs, data analysis, or substantive  
 1921 portions of the manuscript.  
 1922

1923 We assume full responsibility for all content presented in this paper, including any text that has  
 1924 been revised with the assistance of an LLM. We have meticulously reviewed and edited all content  
 1925 to ensure its scientific accuracy and originality, preventing any form of plagiarism or academic  
 1926 misconduct.  
 1927

1928  
 1929  
 1930  
 1931  
 1932  
 1933  
 1934  
 1935  
 1936  
 1937  
 1938  
 1939  
 1940  
 1941  
 1942  
 1943

1944  
 1945  
 1946  
 1947  
 1948  
 1949  
 1950  
 1951  
 1952  
 1953  
 1954  
 1955  
 1956  
 1957  
 1958  
 1959  
 1960  
 1961  
 1962  
 1963  
 1964  
 1965  
 1966  
 1967  
 1968  
 1969  
 1970  
 1971  
 1972  
 1973  
 1974  
 1975  
 1976  
 1977  
 1978  
 1979  
 1980  
 1981  
 1982  
 1983  
 1984  
 1985  
 1986  
 1987  
 1988  
 1989  
 1990  
 1991  
 1992  
 1993  
 1994  
 1995  
 1996  
 1997

Table 40: Training settings for RL.

Parameter	Value
n_gpu	8
rollout.n	16
total_steps	1000
batch_size	8
critic_warmup	0
max_prompt_length	512
max_response_length	16384
filter_overlong_prompts	True
learning_rate	1e-6
use_kl_loss	True
kl_loss_coef	0.001
kl_loss_type	low_var_kl