

Meta-Reasoner: Dynamic Guidance for Optimized Inference-time Reasoning in Large Language Models

Yuan Sui^{1*} Yufei He¹ Tri Cao¹ Simeng Han² Yulin Chen¹ Bryan Hooi¹

¹National University of Singapore ²Yale University

Abstract

Large Language Models (LLMs) often struggle with computational efficiency and error propagation in multi-step reasoning tasks. While recent advancements on prompting and post-training have enabled LLMs to perform step-wise reasoning, they still tend to explore unproductive solution paths without effective backtracking or strategy adjustment. In this paper, we propose Meta-Reasoner, a new framework that empowers LLMs to “think about how to think”. It optimizes the inference process by dynamically adapting reasoning strategies in real-time. Our approach employs contextual multi-armed bandits (CMABs) to learn an adaptive policy. It learns to evaluate the current state of LLM’s reasoning and determine optimal strategy that is most likely to lead to a successful outcome during inference, like whether to backtrack, switch to a new approach, or restart the problem-solving process. This meta-guidance helps avoid unproductive paths exploration during inference and hence improves computational efficiency. We evaluate Meta-Reasoner on math problems (e.g., Game-of-24, TheoremQA) and scientific tasks (e.g., SciBench). Results show that our method outperform previous SOTA methods by 9-12% in accuracy, while reducing inference time by 28-35% under the same compute budget. Additional experiments on creative writing demonstrate the generalizability of our approach to diverse reasoning-intensive tasks. ¹

1 Introduction

Recent advancements on post-training of large language models (LLMs) like o1/o3/r1 (OpenAI et al., 2024; DeepSeek-AI et al., 2025), have achieved remarkable performance on complex reasoning tasks, such as math (Patel et al., 2024; Lightman et al.,

2023), science (Rein et al., 2024), and logical puzzles (Lei et al., 2024; Yao et al., 2023). By simulating human-like deliberation (Yao et al., 2024; Wei et al., 2022), these approaches enable LLMs to decompose problems into subproblems, test hypotheses, reflect on intermediate results, and iteratively refine candidate solutions at inference time (Cao et al., 2024). This extended inference-time reasoning allows models to progressively improve intermediate solutions before committing to a final answer (Chenghao Yang, 2024).

Despite these advances, this deliberate inference-time reasoning remains fundamentally trial-and-error. While this facilitates exploration of diverse solution strategies, recent works on scaling test-time compute show that naively increasing the amount of reasoning (e.g., the number of generated tokens or steps) leads to diminishing returns, where accuracy plateaus and stops improving despite the extra computational cost (Snell et al., 2024; Manvi et al., 2025). Specifically, current approaches face two key limitations: (1) **computational inefficiency**, where models expend substantial inference-time compute on unproductive or redundant reasoning trajectories (Chenghao Yang, 2024); and (2) **error propagation**, where early mistaken assumptions propagate through long chains of reasoning and are difficult to revoke (Lei et al., 2024; Ling et al., 2023). Empirical studies indicate that longer chain-of-thought (CoT) traces can improve accuracy, but they are increasingly expensive and prone to unproductive loops in the absence of higher-level guidance (Havrilla et al., 2024).

Recent techniques such as backtracking and self-reflection can partially mitigate these issues by revising or critiquing model outputs (Gandhi et al., 2024; Li et al., 2025). However, these methods still lack a principled mechanism to *holistically* revise or redirect the ongoing reasoning process. In particular, they provide limited support for dynamically changing the overall strategy (e.g., choos-

* Corresponding Email: yuan.sui@u.nus.edu

¹Code are released at <https://github.com/Y-Sui/Meta-reasoner>

ing to abandon an unpromising decomposition or restart from alternative premises) when the current trajectory is unlikely to succeed (Gao et al., 2024). Additionally, as LLMs are tackling more challenging problems which require longer reasoning, this absence of adaptive, higher-level control becomes more critical. A **critical challenge**, therefore, is to enable LLMs to manage their reasoning budget more effectively, i.e., prioritizing promising directions while adapting or discarding ineffective strategies during inference time. Addressing this requires a novel approach to provide adaptive oversight of the reasoning process.

In this paper, we propose Meta-Reasoner, a meta-reasoning module that operates alongside an LLM to dynamically optimize its reasoning strategy during inference. Acting as a high-level controller, Meta-Reasoner continuously evaluates the LLM’s current reasoning state and provides strategic guidance. Inspired by dual-process theory (Didolkar et al., 2024), we explicitly decouple *high-level strategy selection* (System-2-like) from *low-level stepwise generation* (System-1-like) via a lightweight *progress report* interface. In specific, Meta-Reasoner only considers a compact summary of the LLM’s recent reasoning and then proposes updated strategies based on this summary. The system operates iteratively: (1) the base LLM generates partial CoT reasoning steps (§3.1) and summarizes its current reasoning state into a lightweight progress report (§3.2). (2) Meta-Reasoner reviews this report and uses a contextual multi-armed bandit (CMAB) algorithm (§3.3) to select high-level actions (e.g., backtrack to a previous step, change the decomposition scheme, or restart from alternative axioms); (3) the LLM then continues reasoning, conditioned on both the original context and the controller’s guidance. This design allows Meta-Reasoner to focus on strategic control while leaving fine-grained token-level generation to the base model, thereby reducing overhead and improving robustness. Based on our experiments results (§4), we find that Meta-Reasoner consistently improves performance over strong inference-time reasoning baselines on complex mathematical, logical, and puzzle benchmarks, and that its meta-reasoning strategies transfer to other domains such as creative writing. **Overall, our main contributions are:**

- We introduce a meta-reasoning framework that provides high-level strategic control over LLM inference, reducing the tendency of models to get




Figure 1: **Dynamic Strategy Optimization with CMAB.** It shows how a CMAB algorithm learns to choose the best strategy. It starts with an initial probability distribution for each strategy. The sample process then selects a strategy α_i , and then uses the resulting success or failure feedback r_{t_b} to update both the estimated value (Q-value) of each strategy and the future selection probabilities π_{t_b} . This process runs iteratively to optimize the strategy selection over time.

trapped in unproductive reasoning trajectories.

- We design a lightweight progress reporting mechanism that enables efficient communication of the LLM’s internal reasoning state to the meta-reasoner with minimal inference-time compute.
- We demonstrate that Meta-Reasoner requires no task-specific fine-tuning and generalizes across domains, including mathematics, science, puzzles, and open-ended creative problems.

2 Preliminary

In complex reasoning tasks, a **key challenge** is deciding the best strategy from several valid options. Consider solving a complex math problem: researchers may employ strategies like decomposition, abstraction, heuristic validation, boundary testing or any other methods at any given step. **The critical question is *which* strategy to use and *when*.** This decision making problem aligns well with *contextual multi-armed bandits* (CMABs) problem (Slivkins, 2019), a well-studied framework for making such choices. The framework is designed to balance *exploration* (trying new strategies to learn) and *exploitation* (using strategies known to be effective) based on the current state.

Intuitive Understanding of CMAB. Imagine an agent faced with several strategies (called *arms*), and at each step, it observes some information about the current situation (called the *context*). Based on this context, the agent picks one strategy to apply and then receives feedback (a *reward*) indicating how well that strategy performed. The goal of the CMAB problem is to choose strategies over time to maximize the total reward.

As shown in Figure 1, it illustrates a dynamic strategy selection process using CMAB approach.

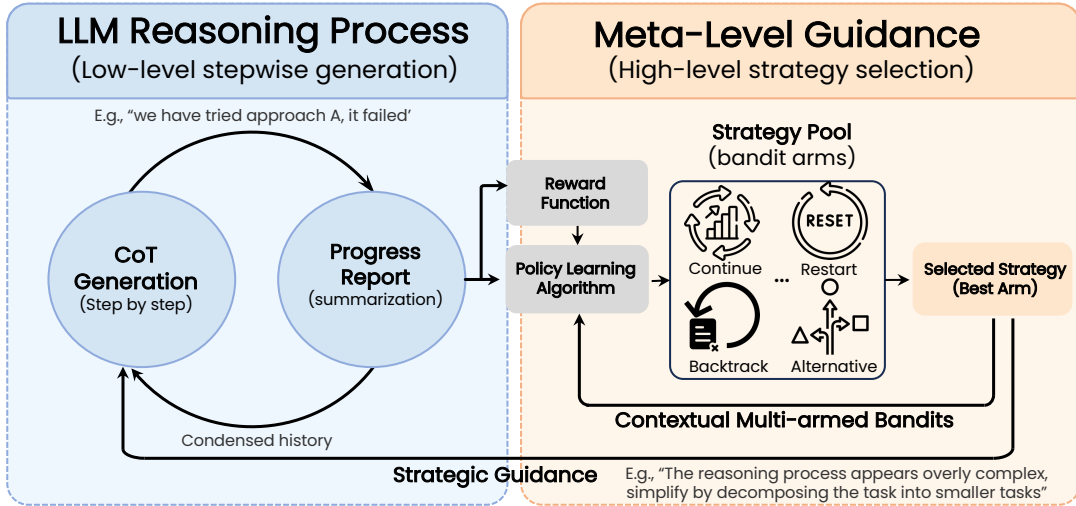


Figure 2: **Workflow of Meta-Reasoner.** In each round, LLM produces a new reasoning step to extend its reasoning trajectories (§3.1). The reasoning process is then summarized into a progress report, which provides context for the meta-reasoner (§3.2). Then meta-reasoner employs a CMAB approach to choose a guidance strategy (§3.3). The selected strategy then guides the next reasoning step generation, to enable strategic redirection, error correction, or resource optimization. A reward is then computed from the progress report and used to update the bandit algorithm (§3.3.1). The process repeats until the task is complete or the maximum number of rounds is reached.

The agent manages three arms which correspond to three strategies, and each with an associated strategy value (Q-value) that reflects its effectiveness. Based on feedback (r_{t_b}) from problem-solving attempts, the agent updates the probability distribution π_{t_b} , which determines the likelihood of selecting each strategy (e.g., strategy 1 might have a 50% chance, strategy 2 a 30% chance, and strategy 3 a 20% chance). The sample process then picks a specific strategy α_i (illustrated as the “chosen strategy”) from this distribution to apply to the current problem. The resulting performance feedback is fed back into the controller, which adjusts the Q-value and refines π_{t_b} for the next iteration, optimizing the strategy selection over time.

Formalization of CMABs. Formally, at each time step t , the model observes a context vector x_t describing the current state and selects an arm s_t from a set of possible strategies \mathcal{S} . Then, the model receives a reward $r(s_t, x_t)$ that depends on both the chosen strategy and the context. The model aims to maximize the cumulative reward over T steps:

$$R(T) = \sum_{t=1}^T r(s_t, x_t). \quad (1)$$

A popular algorithm for this problem is LinUCB (Li et al., 2010), which assumes the expected reward is approximately a linear function of the context. Specifically, for each arm s , there is an unknown parameter vector θ_s such that:

$$\mathbb{E}[r(s, x_t)] \approx x_t^\top \theta_s. \quad (2)$$

LinUCB maintains an estimate $\hat{\theta}_s$ of this parameter and a measure of uncertainty about the estimate. The key idea of LinUCB is its approach to the exploration-exploitation dilemma through an *uncertainty bonus*. It quantifies the statistical uncertainty in the reward estimate: arms with higher estimated uncertainty receive a bonus, ensuring they are explored sufficiently. At each step, LinUCB selects the arm that maximizes a combination of the estimated reward and uncertainty bonus:

$$s_t = \arg \max_{s \in \mathcal{S}} \left[x_t^\top \hat{\theta}_s + c \sqrt{x_t^\top A_s^{-1} x_t} \right], \quad (3)$$

where A_s is a matrix capturing past observations for arm s , and c controls how much the model favors exploring uncertain options. The second term encourages trying arms that might perform well but have been less explored.

By using context to guide its choices, LinUCB adapts its strategy selection to the current situation, which aligns with our goal in Meta-Reasoner: to dynamically choose the most effective reasoning strategy during inference time. We provide the details about how to leverage these CMABs approaches in our system in Section 3.3.

3 Methods

Motivated by the intuition that LLMs should concentrate their computational efforts on more promising reasoning paths during inference, we explore two key research questions in this paper:

- (1) *How can language models dynamically allocate resources during inference to optimize rea-*

Diagnosis	Strategy
Progress is insufficient or the current strategy seems ineffective.	Restart from scratch and propose alternative strategies.
There are mistakes in intermediate steps.	Backtrack to the point where the error occurred.
The current approach is working well.	Continue and provide specific suggestions for the next steps.
Ambiguous or conflicting intermediate results are observed.	Pause to clarify and disambiguate the current reasoning, then reconcile the discrepancies.
The reasoning process appears overly complex or convoluted.	Simplify by decomposing the task into smaller, manageable sub-tasks.
Evidence of error propagation or low confidence in certain sub-components.	Perform targeted verification on critical steps and focus on areas with low confidence.
Repetitive or circular reasoning patterns are detected.	Reset to a previously successful checkpoint and explore alternative solution paths.

Table 1: **Demonstration** of contextual bandit pairs. The top three rows represent standard strategies, while the **highlighted bottom four rows** depict unique strategies generated by Dynamic Contextual Bandits as described in §3.3.3.

soning and planning?

- (2) *What architectural design enables an effective separation between the reasoning process within the LLM and the meta-level guidance that oversees it?*

To address these questions, we propose **Meta-Reasoner**, which endows LLMs with the ability to “think about how to think”. Our framework supervises the LLM’s reasoning process and dynamically guides the model to focus on more promising reasoning trajectories during inference time. The framework operates iteratively as illustrated in Figure 2. At each round t , the reasoning process comprises three steps: (1) *CoT generation* by the LLM, (2) *Progress Reporting* to summarize the reasoning progress so far (i.e., this is partly for efficiency, and partly to help the meta-reasoner focus on its main goal of “advising” rather than being distracted by the details in the CoT), and (3) *Strategy Generation* by the meta-reasoner to optimize subsequent steps. The selection of the strategy almost exactly corresponds to the well-studied problem of contextual multi-armed bandits as illustrated in §2. Each strategy can be seen as an arm for the bandit, and the reward of each strategy can be evaluated by the progress of LLM reasoning after applying the strategy. We analogize the process of executing and evaluating each strategy as the act of “pulling” each arm. The overall goal of our meta-reasoner is to find the best arm (i.e., strategy with highest cumulative rewards) with as few pulls as possible. The complete process of Meta-Reasoner is demonstrated in Algorithm 1. The prompt for each step is detailed in Appendix §F.

3.1 Step-wise CoT Generation

In the first step, the LLM generates a reasoning step s_t to extend its reasoning trajectory $C_t =$

$C_{t-1} \cup \{s_t\}$ based on the user query. This step builds on the previous reasoning history C_{t-1} and incorporates strategic guidance G_{t-1} from the meta-reasoner in the prior round. By maintaining the complete reasoning trajectory step-by-step, the model maintains a coherent foundation for refining the progress of reasoning process. This step resembles the long-term reasoning process observed in models like o1 and r1, which generate extended CoTs during inference. The key difference is that Meta-Reasoner actively intervenes in the reasoning process by providing strategical guidance G_{t-1} , enabling the LLMs “think about how to think”.

3.2 Progress Reporting

After updating the CoT, we apply a summarization function $f : C_t \rightarrow P_t$ to distill the most recent t steps of the CoT into a concise progress report P_t . This report encapsulates critical elements of the reasoning trajectory, including the degree of progress toward the task objective, the logical consistency of the reasoning steps, and significant milestones or updates achieved. The function $f(\cdot)$ is designed to be computationally efficient, producing a compact summary that enables the meta-reasoner to assess **high-level** progress without processing the **granular details** of each step in C_t . While this summarization is heuristic in nature, we observe that it may trigger LLMs’ capacity for higher-order reasoning by including only essential information in P_t , which informs the meta-reasoner’s strategic guidance G_t . Empirically, this approach fosters more insightful and adaptive strategies, as the meta-reasoner can focus on critical patterns rather than exhaustive step-by-step details. We provide the detailed experiments in §4.2.

Algorithm 1 Meta-Reasoning with Contextual Multi-Armed Bandits

Require: LLM M , Bandits \mathcal{B} , Initial strategy set \mathcal{A}_1 , Maximum rounds T

Ensure: Final answer A_{final}

```
1:  $C_0 \leftarrow \emptyset$ ;  $\mathcal{B}.\text{Initialize}(\mathcal{A}_1)$ 
2:  $G_0 \leftarrow$  default strategy
3: for  $t = 1$  to  $T$  do
4:   if  $t > 1$  then
5:      $P_{t-1} \leftarrow f(C_{t-1})$ 
6:      $x_{t-1} \leftarrow \text{FeatureExtract}(P_{t-1})$ 
7:     (Optional):  $\mathcal{A}_t \leftarrow \mathcal{A}_{t-1} \cup \{\text{new strategies}\}$ 
8:      $a_{t-1} \leftarrow \arg \max_{a \in \mathcal{A}_t} \text{Score}_{\mathcal{B}}(x_{t-1}, a)$ 
9:      $G_t \leftarrow a_{t-1}$ 
10:  else
11:     $G_t \leftarrow G_0$ 
12:  end if
13:   $s_t \leftarrow M(C_{t-1}, G_t)$ 
14:   $C_t \leftarrow C_{t-1} \cup \{s_t\}$ 
15:   $r_t \leftarrow \text{ComputeReward}(C_t)$ 
16:  if  $t > 1$  then
17:     $\mathcal{B}.\text{Update}(x_{t-1}, a_{t-1}, r_t)$ 
18:  end if
19:  if termination condition met then
20:    break
21:  end if
22: end for
23:  $A_{\text{final}} \leftarrow \text{ExtractAnswer}(C_t)$ 
24: return  $A_{\text{final}}$ 
```

3.3 Meta-reasoner Strategy Generation

In the next step, the meta-reasoner evaluates the progress report P_t and selects an appropriate strategy G_t for LLM reasoning (the complete procedure is detailed in Algorithm 1). We formulate the generation of strategy as a CMAB problem (as defined in §2) and explore two settings: (1) a *fixed-strategy* formulation, where the meta-reasoner selects from a predefined set of strategies using a contextual bandit algorithm; and (2) an *advanced* setting, where the meta-reasoner, implemented as an LLM-based agent, dynamically generates or refines strategies. In both settings, the meta-reasoner leverages the partial-feedback mechanism of MABs to adaptively select strategies based on a reward function that evaluates the quality of reasoning progress after applying the chosen strategy G_t . We demonstrate the contextual bandit pair (diagnosis of the progress report (i.e., context) and the corresponding strategy (i.e., bandit) in Table 1.

3.3.1 Reward Modeling for Progress Report.

The primary goal of our evaluation mechanism is to quantify how effectively the model’s current reasoning advances toward the task objective (e.g., solving a complex problem), while also monitoring computational cost to promote efficiency. We define a reward function $R : R : P_t \times G_t \rightarrow \mathbb{R}$ that integrates two key components: (1) **solution**

progress S_p , measuring correctness and adherence to query’s constraints (e.g., $S_p = w_1 \cdot C_c + w_2 \cdot C_a$, where C_c is correctness, C_a is adherence, and w_1, w_2 are weights); and (2) **resource usage** R_u , reflecting computational cost (e.g., $R_u = -\alpha \cdot N_s$, where N_s is the number of reasoning steps and α is a cost coefficient). The total reward is computed as $R(G_t, P_t) = \beta \cdot S_p + (1 - \beta) \cdot R_u$, with $\beta \in [0, 1]$ balancing the trade-off. This evaluation can be implemented via an LLM-based evaluators or external scoring scripts, producing a cumulative score to update the CMAB algorithm. To mitigate potential self-favor bias in LLM-based evaluation (Li et al., 2024; Gu et al., 2024), where a model may favor its own outputs. We use different models for solution generation and reward modeling. Unless otherwise specified, we use gemini-2.5-flash as the default reward modeling model. To address potential concerns regarding the choice of hyperparameters in the reward function, we conduct a sensitivity analysis in Appendix §C.

3.3.2 Fixed Contextual Bandit.

In the basic version of our framework, the meta-reasoner is modeled as a single contextual bandit that selects from a *fixed, finite* set of K strategies. These strategies include instructions such as “continue and provide specific suggestions”, “restart from scratch”, “backtrack to the point where the error occurred”, or “propose alternative methods or perspectives to consider”, as detailed in Table 1. At each round, the LLM produces a *progress report* P_t summarizing its last several reasoning steps, the meta-reasoner then transforms this progress report into a feature vector x_t using a language model and applies a contextual bandit algorithm (e.g., LinUCB (Li et al., 2010)) to select the best next strategy a_t . The LLM then executes that strategy and we collect the reward r_t for a_t based on the reward function. Through iterative MAB algorithm updating, the MAB algorithm learns to select appropriate strategies conditioned on the recent progress report.

3.3.3 Dynamic Contextual Bandit.

The fixed-arm formulation assumes a static set of strategies G_0 . In practice, the meta-reasoner may itself be an LLM capable of inventing new strategies over time. To accommodate *dynamic* strategies, we allow the meta-reasoner to propose or refine new strategies at round t , which generates an expanding collection of strategies, $G_1 \subseteq G_2 \subseteq \dots \subseteq G_t$. Each newly introduced strategy becomes an addi-

Method	Accuracy (%)
GPT-4o-mini + CoT (Yao et al., 2023)	4
GPT-4o-mini + SC-CoT (Yao et al., 2023)	9
GPT-4o-mini + IO (best of 100) (Yao et al., 2023)	33
GPT-4o-mini + CoT (best of 100) (Yao et al., 2023)	49
Gemini-Exp-1206 + IO (best of 100) (Yao et al., 2023)	38
Gemini-Exp-1206 + CoT (best of 100) (Yao et al., 2023)	60

GPT-4o-mini + ToT ($b = 1$) (Yao et al., 2023)	32
GPT-4o-mini+ ToT ($b = 5$) (Yao et al., 2023)	65
GPT-4o-mini + Reflexion (Shinn et al., 2023)	53
GPT-4o-mini + MACM (Lei et al., 2024)	80
GPT-4o-mini + Meta-Reasoner (our work)	89
GPT-4o + Meta-Reasoner (our work)	92
Gemini-Exp-1206 + Meta-Reasoner (our work)	94
Qwen3-8B + Meta-Reasoner (our work)	87
DS-R1-Distill-Qwen-14B + Meta-Reasoner (our work)	93

o1-mini + IO	89
o1-preview + IO	93

Table 2: Performance of different models on 24-points Game (Yao et al., 2023) (b : search breadth)

Method	Accuracy (%)
GPT-4o-mini + CoT	39.46
Gemini-Exp-1206 + CoT	43.12

GPT-4o-mini + Reflexion (Shinn et al., 2023)	74.32
GPT-4 Turbo + MACM (Lei et al., 2024)	79.41
GPT-4o-mini + Meta-Reasoner (our work)	84.13
Gemini-Exp-1206 + Meta-Reasoner (our work)	86.32
Qwen3-8B + Meta-Reasoner (our work)	82.93
DS-R1-Distill-Qwen-14B + Meta-Reasoner (our work)	87.40

Table 3: Performance of different models on TheoremQA (Chen et al., 2023)

tional arm in the contextual multi-armed bandit framework. To encourage at least some exploration on this new arm, we initialize their bandit parameters with a neutral prior, such as $Q(a) = 0$ for arm a to avoid strong biases. We further analyze the stability of these new generated dynamic contextual bandits in Appendix §D, ensuring that the system avoids incorporating random or ineffective strategies by grounding new arms in contextual relevance and iterative feedback.

By explicitly separating low-level content generation (handled by the LLM) from high-level strategy decisions (governed by the meta-reasoner’s bandit), the system can effectively avoid getting stuck or wasting excessive resources on poor reasoning paths. In domains where a predefined set of strategies is sufficient, the fixed-arm formulation can simplify the method deployment. While in more open-ended domains where novel tactics may emerge, dynamic-arm extensions give meta-reasoner more flexibility to evolve.

4 Experiments

Due to page limits, we put all the experiment settings and training details in Appendix §B. In this

Method	Diff(%)	Stat(%)	Calc(%)
Phi-4 + CoT	17.42	28.42	32.93
Llama-3.1-instruct + CoT	33.14	49.72	54.18
Gemini-Exp-1206 + CoT	36.32	56.73	59.24
Gemini-Exp-1206 + SC-CoT	38.73	59.12	64.11
GPT-4o-mini + CoT	33.12	55.71	58.10
GPT-4o-mini + SC-CoT	37.33	56.67	63.81
GPT-4o-mini + MCR	40.12	58.21	67.42

GPT-4o-mini + MACM (Lei et al., 2024)	54.78	67.13	65.77
GPT-4o + MACM (Lei et al., 2024)	61.42	78.32	76.72
GPT-4o-mini + Meta-Reasoner (our work)	60.32	73.64	80.23
GPT-4o + Meta-Reasoner (our work)	67.14	83.29	84.17
Qwen3-8B + Meta-Reasoner (our work)	59.31	71.80	78.65
DS-R1-Distill-Qwen-14B + Meta-Reasoner (our work)	66.70	78.43	84.55

Table 4: Performance of different models on SciBench Math Subsets (Wang et al., 2023)

section, we present the main results, ablation study, analysis regarding efficiency, rewards accumulation, and qualitative assessment.

4.1 Main Results

We compare the accuracy of different prompting methods across different backbone models on SciBench (as shown in Table 4), 24-points game (as shown in Table 2) and TheoremQA (as shown in Table 3). We find that basic prompting strategies, such as CoT and SC-CoT, show limited effectiveness, achieving only 4% and 9% accuracy on 24-point games, respectively. Incorporating IO strategy with “Best of 100” samples improves accuracy to 33%, but it remains far behind advanced methods like MACM or ToT. Strategies like ToT illustrate the performance of exploring broader reasoning paths in a structured manner, with accuracy increasing from 45% to 74%. The more advanced iterative methods like Reflexion and MACM further demonstrate the value of refined reasoning frameworks incorporating multi-step reflection and error checking. Our proposed Meta-Reasoner outperforms these approaches, achieving 89% accuracy with GPT-4o-mini and 92% with GPT-4o, showcasing its ability to dynamically guide reasoning, correct errors, and focus resources effectively. Compared to specialized models like o1-mini, our method equipped with much cheaper models like GPT-4o-mini delivers comparable performance.

Overall, the Meta-Reasoner framework provides a compatible approach to reasoning-intensive tasks, achieving high accuracy with dynamic and efficient problem-solving strategies. The results on SciBench (Table 4) and TheoremQA (Table 3) also demonstrate similar findings and show that Meta-Reasoner generally achieves better performance compared to the baselines and the results are consistent across different models.

Method	Base Model	MATH-500 (%)	AIME 2024 (Pass@1)
rStar-Math (Guan et al., 2025)	Qwen2.5-Math-7B	90.0	53.3
rStar-Math (Guan et al., 2025)	Phi3-mini-3.8B	86.4	43.3
MCTS-RAP (Hao et al., 2023)	GPT-4o-mini	72.5	14.5
MCTS-RAP (Hao et al., 2023)	GPT-4o	79.0	23.4
MCTS-RAP (Hao et al., 2023)	DeepSeek-R1-Distill-Qwen-7B	84.3	40.2
Meta-Reasoner (Ours)	GPT-4o-mini	85.6	26.7
Meta-Reasoner (Ours)	GPT-4o	87.3	33.3
Meta-Reasoner (Ours)	DeepSeek-R1-Distill-Qwen-7B	92.3	55.5

Table 5: Performance on Math-500 (Lightman et al., 2023) and AIME-2024 (Guan et al., 2025; AI-MO Team, 2024).

Model	Variant	Game-of-24	TheoremQA
GPT-4o-mini	Full Method	89	84.13
	w/o Progress Report	85	79.42
	w/o MAB (direct arm selection)	82	80.74
	w/o MAB (CoT)	4	39.46
Gemini-Exp-1206	Full Method	94	86.32
	w/o Progress Report	91	81.78
	w/o MAB (direct arm selection)	87	82.14
	w/o MAB (CoT)	11	43.12

Table 6: Ablation study of Meta-Reasoner. Direct arm selection refers to prompting LLM to directly select a strategy based on recent progress report.

4.2 Ablation Study

We further conduct an ablation study to analyze each component contribution of Meta-Reasoner. Specifically, we consider the following setup: (1) w/o progress report: we replace the progress reporting process with directly considering the entire CoT history without summarization; (2) w/o MAB: instead of using MAB to select the proper strategy, we directly prompting an LLM to provide the reasoning strategy.

In Table 6, we show that when removing progress reporting (“w/o Progress Report”), the overall performance moderately degrades and we hypothesize it is due to the concise intermediate summarizations can help Meta-Reasoner only consider the high-level strategy instead of being confused with too many details of the reasoning process. We also find that removing the MAB brings a more pronounced effect, especially when strategy selection falls back to a direct chain-of-thought approach (“w/o MAB (CoT)”). It verifies the effect of our meta-reasoner module to help the model stay on track for getting an optimal solution. In Table 7, we compare fixed and dynamic bandit variants on the game of 24 and theoremQA. We find that using a fixed set of strategies (e.g., $K = 3$ and $K = 5$) yields lower performance compared to the dynamic approach which adaptively explores more strategies (shown by larger unique strategies). The results highlight the benefit of flexibly allocating diverse reasoning strategies using LLM in-context learning capabilities.

4.3 Analysis

Effectiveness of Meta-reasoner. Figure 3 demonstrates the cumulative rewards across iterations. We compare our MAB-based approach with a baseline that directly prompts an LLM to select an arm (or “strategy”), referred to as *Baseline (Direct Arm Selection)*; the prompt details are in Appendix §F. Results show that the MAB-based meta-reasoner (using LinUCB (Li et al., 2010)) consistently outperforms both direct LLM decision-making and random search across two tasks (Game of 24 and TheoremQA) and two model scales (GPT-4o-mini and Gemini-Exp-1206). While direct LLM prompting yields reasonable initial performance and random search requires minimal setup, neither approach effectively balances exploration and exploitation. In contrast, the MAB updating strategy leverages feedback from prior iterations to adaptively refine action selection (e.g., choosing an appropriate strategy based on CoT reasoning), steadily increasing cumulative rewards.

Bandit Type	Game-of-24(%)	#US	TheoremQA(%)	#US
Fixed (K=3)	65	3	72.34	3
Fixed (K=5)	72	5	79.17	5
Dynamic	89	14	84.13	21

Table 7: Fixed vs. Dynamic Bandit Variants over GPT-4o-mini. #US: Number of Unique Strategies.

Inference Efficiency. We measure the inference efficiency of our proposed method. In Figure 4, we calculate the average inference time on different tasks across different models and methods. It demonstrates Meta-Reasoner’s significant advantages over existing reasoning frameworks across multiple benchmarks. Compared to o1-preview in the zero-shot setting, our method achieved roughly 51-55% inference time reduction while maintaining comparable performance. Among reasoning methods using identical base models, Meta-Reasoner consistently outperformed alternatives: requiring 29-33% less inference time than ToT (Yao et al., 2023), Reflexion (Shinn et al., 2023), MACM (Lei et al., 2024) and Best-of-N techniques. We further provide the demonstration of the token usage comparison in Figure 5. The efficiency results positions Meta-Reasoner as a scalable solution for complex reasoning tasks where computational resources are constrained.

Iterative Reasoning and Memory Management. Unlike traditional CoT methods that generate rea-

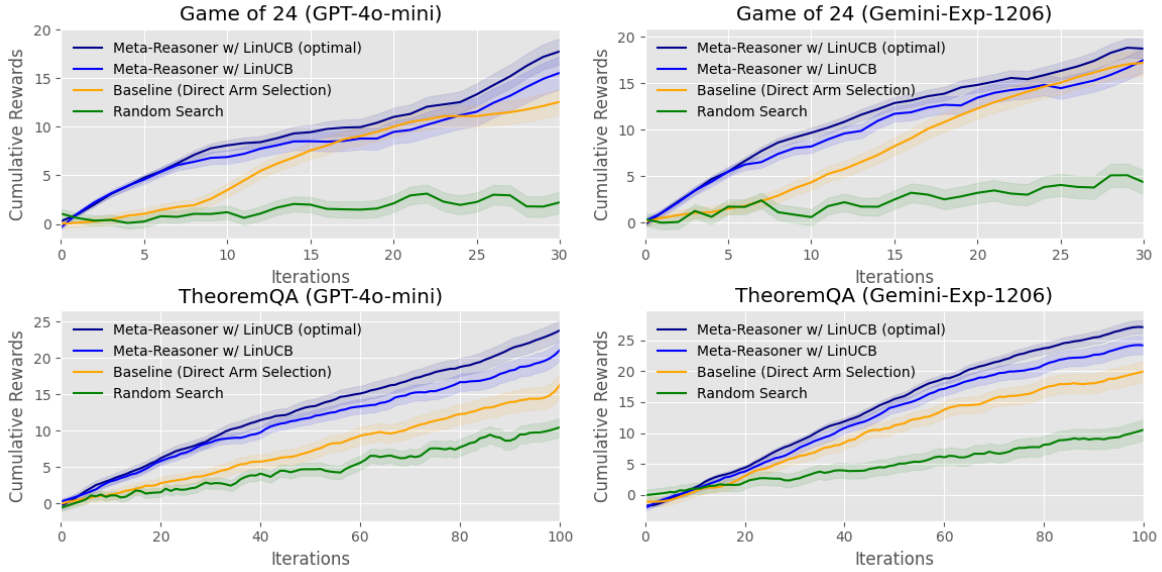



Figure 3: Cumulative reward of different settings across iteration. We compare our method using LinUCB with baseline (direct arm selection), and random search methods across two tasks—Game of 24 (top row) and TheoremQA (bottom row) using GPT-4o-mini (left) and Gemini-Exp-1206 (right).



(a) Game of 24 Task. (b) TheoremQA Task
Figure 4: Inference time heatmap comparison.

soning process in a single pass, Meta-Reasoner pauses generation periodically to evaluate the reasoning progress and adaptively switch the reasoning strategies using a multi-armed bandit-based algorithm. This iterative design adds additional memory overhead due to maintaining context across multiple LLM calls. To mitigate this, we summarize the entire reasoning history into compact progress reports (§3.2), greatly reducing the token load per iteration. As demonstrated in Figure 4, the savings achieved by avoiding unproductive reasoning paths significantly outweigh the modest overhead introduced by the strategy selection process.

Generalizability of Meta-Reasoner. We test the generalizability of Meta-Reasoner on creative writing task, which differs from math and logical reasoning. Following ToT (Yao et al., 2023), the input consists of four random sentences, and the model must generate a four-paragraph story that ends with those sentences. We use the same automatic evaluation protocol as ToT: a GPT-4 zero-shot judge assigns a scalar coherence score from 1 to 10 for

each output. Table 8 shows that Meta-Reasoner achieves a coherence score of 7.68. This is competitive with ToT + Prompt Refined (7.91) and clearly better than IO (6.19) and CoT (6.93). These results suggest that Meta-Reasoner is domain-agnostic, it can transfer beyond mathematical objectives and remains effective in open-ended creative generation.

Method	GPT-4 Coherence Score
IO (Input-Output)	6.19
IO + Prompt Refined ($k = 5$)	7.67
CoT (Chain-of-Thought) (Wei et al., 2022)	6.93

ToT (Tree-of-Thoughts) (Yao et al., 2023)	7.56
ToT + Prompt Refined (Yao et al., 2023)	7.91
Meta-Reasoner	7.68

Table 8: Generalizability of Meta-Reasoner on the creative writing task (Yao et al., 2023). Higher GPT-4 coherence scores are better.

5 Conclusion

In this work, we introduce Meta-Reasoner, a framework designed to enhance the reasoning capabilities of LLMs and optimize the inference-time efficiency. By operating as an “advisor”, meta-reasoner dynamically evaluates the reasoning process and provides high-level strategic guidance, addressing key limitations of long CoT reasoning, such as compounding errors and inefficiency in inference computing. Unlike conventional approaches, Meta-Reasoner focuses on global oversight rather than granular step-by-step reasoning, enabling LLMs to avoid unproductive lines

of thought and better allocate computational resources. Experiments highlight the potential of dynamic reasoning to overcome inherent challenges in the LLM reasoning and also show promise in broader applications, offering a scalable and adaptable solution for reasoning-intensive tasks.

Limitations

Our proposed Meta-Reasoner framework, while effective at enhancing inference-time reasoning, remains limited to text-based problems and struggles to address tasks requiring other modalities, such as geometry. Overcoming these challenges calls for further advancements in the model’s cognitive capabilities.

References

- AI-MO Team. 2024. Aimo validation aime dataset. <https://huggingface.co/datasets/AI-MO/aimo-validation-aime>. Dataset containing problems from AIME 2022-2024.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Michal Podstawski, Hubert Niewiadomski, Piotr Nyczyk, et al. 2023. [Graph of thoughts: Solving elaborate problems with large language models](#). *arXiv preprint arXiv:2308.09687*.
- Yuji Cao, Huan Zhao, Yuheng Cheng, Ting Shu, Guolong Liu, Gaoqi Liang, Junhua Zhao, and Yun Li. 2024. [Survey on Large Language Model-Enhanced Reinforcement Learning: Concept, Taxonomy, and Methods](#). *arXiv preprint*.
- Wenhu Chen, Ming Yin, Max Ku, Pan Lu, Yixin Wan, Xueguang Ma, Jianyu Xu, Xinyi Wang, and Tony Xia. 2023. [TheoremQA: A theorem-driven question answering dataset](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7889–7901, Singapore. Association for Computational Linguistics.
- Chenghao Yang. 2024. [Inference Time Compute](#).
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanbiao Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yuxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv: 2501.12948*.
- Aniket Rajiv Didolkar, Anirudh Goyal, Nan Rosemary Ke, Siyuan Guo, Michal Valko, Timothy P Lillicrap, Danilo Jimenez Rezende, Yoshua Bengio, Michael Curtis Mozer, and Sanjeev Arora. 2024. [Metacognitive capabilities of LLMs: An exploration in mathematical problem solving](#). In *AI for Math Workshop @ ICML 2024*.
- Kanishk Gandhi, Denise H. J. Lee, Gabriel Grand, Muxin Liu, Winson Cheng, Archit Sharma, and Noah Goodman. 2024. [Stream of Search \(SoS\): Learning to Search in Language](#). In *First Conference on Language Modeling*.
- Peizhong Gao, Ao Xie, Shaoguang Mao, Wenshan Wu, Yan Xia, Haipeng Mi, and Furu Wei. 2024. [Meta reasoning for large language models](#). *arXiv preprint arXiv: 2406.11698*.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Yuanzhuo Wang, and Jian Guo. 2024. [A survey on llm-as-a-judge](#). *CoRR*, abs/2411.15594.

- Xinyu Guan, Li Lyna Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang, and Mao Yang. 2025. [rstar-math: Small llms can master math reasoning with self-evolved deep thinking](#). *arXiv preprint arXiv: 2501.04519*.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujia Yang. 2023. [Evoprompt: Connecting llms with evolutionary algorithms yields powerful prompt optimizers](#). *arXiv preprint arXiv: 2309.08532*.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujia Yang. 2025. [Evoprompt: Connecting llms with evolutionary algorithms yields powerful prompt optimizers](#). *Preprint*, arXiv:2309.08532.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. 2023. Reasoning with language model is planning with world model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8154–8173.
- Alexander Havrilla, Sharath Chandra Raparthy, Christoforos Nalmpantis, Jane Dwivedi-Yu, Maksym Zhuravinskyi, Eric Hambro, and Roberta Raileanu. 2024. [Glore: When, where, and how to improve llm reasoning via global and local refinements](#). In *ICML*.
- Kuang-Huei Lee, Ian Fischer, Yueh-Hua Wu, Dave Marwood, Shumeet Baluja, Dale Schuurmans, and Xinyun Chen. 2025. [Evolving deeper llm thinking](#). *arXiv preprint arXiv: 2501.09891*.
- Bin Lei, Yi Zhang, Shan Zuo, Ali Payani, and Caiwen Ding. 2024. [MACM: Utilizing a multi-agent system for condition mining in solving complex mathematical problems](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. 2024. [Llms-as-judges: A comprehensive survey on llm-based evaluation methods](#). *CoRR*, abs/2412.05579.
- Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. [A contextual-bandit approach to personalized news article recommendation](#). In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, page 661–670, New York, NY, USA. Association for Computing Machinery.
- Xiaoxi Li, Guanting Dong, Jiajie Jin, Yuyao Zhang, Yujia Zhou, Yutao Zhu, Peitian Zhang, and Zhicheng Dou. 2025. [Search-o1: Agentic search-enhanced large reasoning models](#). *arXiv preprint arXiv: 2501.05366*.
- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, I. Sutskever, and K. Cobbe. 2023. [Let’s verify step by step](#). *International Conference on Learning Representations*.
- Zhan Ling, Yunhao Fang, Xuanlin Li, Zhiao Huang, Mingu Lee, Roland Memisevic, and Hao Su. 2023. [Deductive verification of chain-of-thought reasoning](#). In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23*, Red Hook, NY, USA. Curran Associates Inc.
- Rohin Manvi, Anikait Singh, and Stefano Ermon. 2025. [Adaptive inference-time compute: LLMs can predict if they can do better, even mid-generation](#).
- OpenAI, :, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, Ananya Kumar, Andre Saraiva, Andrea Vallone, Andrew Duberstein, Andrew Kondrich, Andrey Mishchenko, Andy Applebaum, Angela Jiang, Ashvin Nair, Barret Zoph, Behrooz Ghorbani, Ben Rossen, Benjamin Sokolowsky, Boaz Barak, Bob McGrew, Borys Minaiev, Botao Hao, Bowen Baker, Brandon Houghton, Brandon McKinzie, Brydon Eastman, Camillo Lugaresi, Cary Bassin, Cary Hudson, Chak Ming Li, Charles de Bourcy, Chelsea Voss, Chen Shen, Chong Zhang, Chris Koch, Chris Orsinger, Christopher Hesse, Claudia Fischer, Clive Chan, Dan Roberts, Daniel Kappler, Daniel Levy, Daniel Selsam, David Dohan, David Farhi, David Mely, David Robinson, Dimitris Tsipras, Doug Li, Dragos Oprica, Eben Freeman, Eddie Zhang, Edmund Wong, Elizabeth Proehl, Enoch Cheung, Eric Mitchell, Eric Wallace, Erik Ritter, Evan Mays, Fan Wang, Felipe Petroski Such, Filippo Raso, Florencia Leoni, Foivos Tsimpourlas, Francis Song, Fred von Lohmann, Freddie Sulit, Geoff Salmon, Giambattista Parascandolo, Gildas Chabot, Grace Zhao, Greg Brockman, Guillaume Leclerc, Hadi Salman, Haiming Bao, Hao Sheng, Hart Andrin, Hessam Bagherinezhad, Hongyu Ren, Hunter Lightman, Hyung Won Chung, Ian Kivlichan, Ian O’Connell, Ian Osband, Ignasi Clavera Gilaberte, Ilge Akkaya, Ilya Kostrikov, Ilya Sutskever, Irina Kofman, Jakub Pachocki, James Lennon, Jason Wei, Jean Harb, Jerry Twore, Jiacheng Feng, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joaquin Quiñero Candela, Joe Palermo, Joel Parish, Johannes Heidecke, John Hallman, John Rizzo, Jonathan Gordon, Jonathan Uesato, Jonathan Ward, Joost Huizinga, Julie Wang, Kai Chen, Kai Xiao, Karan Singhal, Karina Nguyen, Karl Cobbe, Katy Shi, Kayla Wood, Kendra Rimbach, Keren Gu-Lemberg, Kevin Liu, Kevin Lu, Kevin Stone, Kevin Yu, Lama Ahmad, Lauren Yang, Leo Liu, Leon Maksin, Leyton Ho, Liam Fedus, Lilian Weng, Linden Li, Lindsay McCallum, Lindsey Held, Lorenz Kuhn, Lukas Kondrasiuk, Lukasz Kaiser, Luke Metz, Madelaine Boyd, Maja Trebacz, Manas Joglekar, Mark Chen, Marko Tintor, Mason Meyer, Matt Jones, Matt Kaufer, Max Schwarzer, Meghan Shah, Mehmet Yatbaz, Melody Y. Guan, Mengyuan Xu, Mengyuan Yan, Mia Glaese, Mianna Chen, Michael Lampe, Michael Malek, Michele Wang, Michelle Fradin, Mike McClay, Mikhail Pavlov, Miles Wang, Mingx-

- uan Wang, Mira Murati, Mo Bavarian, Mostafa Rohaninejad, Nat McAleese, Neil Chowdhury, Neil Chowdhury, Nick Ryder, Nikolas Tezak, Noam Brown, Ofir Nachum, Oleg Boiko, Oleg Murk, Olivia Watkins, Patrick Chao, Paul Ashbourne, Pavel Izmailov, Peter Zhokhov, Rachel Dias, Rahul Arora, Randall Lin, Rapha Gontijo Lopes, Raz Gaon, Reah Miyara, Reimar Leike, Renny Hwang, Rhythm Garg, Robin Brown, Roshan James, Rui Shu, Ryan Cheu, Ryan Greene, Saachi Jain, Sam Altman, Sam Toizer, Sam Toyer, Samuel Miserendino, Sandhini Agarwal, Santiago Hernandez, Sasha Baker, Scott McKinney, Scottie Yan, Shengjia Zhao, Shengli Hu, Shibani Santurkar, Shraman Ray Chaudhuri, Shuyuan Zhang, Siyuan Fu, Spencer Papay, Steph Lin, Suchir Balaji, Suvansh Sanjeev, Szymon Sidor, Tal Broda, Aidan Clark, Tao Wang, Taylor Gordon, Ted Sanders, Tejal Patwardhan, Thibault Sottiaux, Thomas Degry, Thomas Dimson, Tianhao Zheng, Timur Garipov, Tom Stasi, Trapit Bansal, Trevor Creech, Troy Peterson, Tyna Eloundou, Valerie Qi, Vineet Kosaraju, Vinnie Monaco, Vitchyr Pong, Vlad Fomenko, Weiyi Zheng, Wenda Zhou, Wes McCabe, Wojciech Zaremba, Yann Dubois, Yinghai Lu, Yinling Chen, Young Cha, Yu Bai, Yuchen He, Yuchen Zhang, Yunyun Wang, Zheng Shao, and Zhuohan Li. 2024. Openai o1 system card. *arXiv preprint arXiv: 2412.16720*.
- Bhrij Patel, Souradip Chakraborty, Wesley A. Suttle, Mengdi Wang, Amrit Singh Bedi, and Dinesh Manocha. 2024. [AIME: AI system optimization via multiple LLM evaluators](#).
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2024. [GPQA: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling*.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik R Narasimhan, and Shunyu Yao. 2023. [Reflection: language agents with verbal reinforcement learning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Aleksandrs Slivkins. 2019. [Introduction to multi-armed bandits](#). *Found. Trends Mach. Learn.*
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv: 2408.03314*.
- Yuan Sui, He Yufei, Ding Zifeng, and Bryan Hooi. 2025. [Can knowledge graphs make large language models more trustworthy? an empirical study over open-ended question answering](#). In *The 63rd Annual Meeting of the Association for Computational Linguistics*.
- Xiaoxuan Wang, Ziniu Hu, Pan Lu, Yanqiao Zhu, Jieyu Zhang, Satyen Subramaniam, Arjun R. Loomba, Shichang Zhang, Yizhou Sun, and Wei Wang. 2023. [Scibench: Evaluating college-level scientific problem-solving abilities of large language models](#). *International Conference on Machine Learning*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Yixuan Weng, Minjun Zhu, Fei Xia, Bin Li, Shizhu He, Shengping Liu, Bin Sun, Kang Liu, and Jun Zhao. 2023. [Large language models are better reasoners with self-verification](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zengqi Wen, Chonghua Liao, and Jianhua Tao. 2024. Beyond examples: High-level automated reasoning paradigm in in-context learning via mcts. *arXiv preprint arXiv: 2411.18478*.
- Jinyang Wu, Mingkuan Feng, Shuai Zhang, Feihu Che, Zengqi Wen, Chonghua Liao, and Jianhua Tao. 2025. [Beyond examples: High-level automated reasoning paradigm in in-context learning via mcts](#). *Preprint, arXiv:2411.18478*.
- Huanjin Yao, Jiaying Huang, Wenhao Wu, Jingyi Zhang, Yibo Wang, Shunyu Liu, Yingjie Wang, Yuxin Song, Haocheng Feng, Li Shen, and Dacheng Tao. 2024. Mulberry: Empowering mllm with o1-like reasoning and reflection via collective monte carlo tree search. *arXiv preprint arXiv: 2412.18319*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Ori Yoran, Tomer Wolfson, Ben Bogin, Uri Katz, Daniel Deutch, and Jonathan Berant. 2023. [Answering questions by meta-reasoning over multiple chains of thought](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Che Zhang, Zhenyang Xiao, Chengcheng Han, Yixin Lian, and Yuejian Fang. 2024. Learning to check: Unleashing potentials for self-correction in large language models. *arXiv preprint arXiv:2402.13035*.
- Qiyuan Zhang, Fuyuan Lyu, Zexu Sun, Lei Wang, Weixu Zhang, Wenyue Hua, Haolun Wu, Zhihan Guo, Yufei Wang, Niklas Muennighoff, et al. 2025. A survey on test-time scaling in large language models: What, how, where, and how well? *arXiv preprint arXiv:2503.24235*.

A Related Works

We first review the landscape of complex reasoning in LLMs, focusing on the transition from static Chain-of-Thought (CoT) to dynamic inference-time search. We then discuss the dual-process systems from cognitive science to frame the design of our approach Meta-Reasoner.

Complex Reasoning & Inference-Time Compute. Standard Chain-of-Thought (CoT) (Wei et al., 2022) reasoning has been the cornerstone of LLM problem-solving (Lee et al., 2025). Recent advancements, such as OpenAI’s o1/o3 and Deepseek’s r1, demonstrate that scaling test-time computation to allow the model to "think" longer, can yield state-of-the-art performance (Manvi et al., 2025; Li et al., 2025). However, CoT relies heavily on a strict step-by-step sequence. Consequently, a single error in an early step often cascades, causing the entire reasoning chain to fail (Snell et al., 2024; Zhang et al., 2025). Furthermore, when models face difficult tasks, they often get stuck in repetitive loops without realizing they are making no progress (Sui et al., 2025). Meta-Reasoner addresses these issues by introducing a meta-controller that actively monitors progress and intervenes when the reasoning process stalls or goes off track.

Self-Correction and Verification. To address the issues of error propagation, recent works have introduced mechanisms for self-correction and backtracking (Yao et al., 2023; Besta et al., 2023; Gandhi et al., 2024). Methods like Reflexion (Shinn et al., 2023) and self-verification frameworks (Weng et al., 2023; Ling et al., 2023) ask the model to review and revise its own output. While effective, these approaches often rely on fixed heuristics (e.g., "always verify after step N ") or simple scalar scores that only indicate if a step is "good" or "bad" (Lightman et al., 2023; Zhang et al., 2024). Meta-Reasoner improves upon this by using a learned policy using multi-armed bandits to make specific decisions. Instead of just giving a score, our meta-reasoner provides actionable instructions such as "backtrack to the previous step" or "restart with a new method", allowing for more flexible and intelligent error correction than standard verification loops.

Meta-Reasoner vs. Tree Search. Search algorithms like Tree-of-Thoughts (ToT) (Yao et al.,

2023) and MCTS (Wu et al., 2025) treat reasoning as a search problem, trying to find the best possible next step at every point. However, this approach is computationally expensive because it requires managing a vast number of potential steps. In contrast, Meta-Reasoner operates at a higher level: it functions as a *strategy* controller rather than a *step* controller. Instead of micro-managing every word the model generates, our framework periodically evaluates the overall path and selects high-level strategies via Contextual Multi-Armed Bandits. This allows Meta-Reasoner to guide the general direction of the solution without the high computational cost of expanding a full search tree.

Meta-Cognition & Dual-Process Systems. From a cognitive science perspective, meta-cognition involves higher-order processes that allow individuals to monitor, evaluate, and adjust their cognitive strategies (Gao et al., 2024; Yoran et al., 2023) during thinking process. This reflective thinking—often characterized as System 2 in dual-process theories (Havrilla et al., 2024)—is vital for tasks requiring careful deliberation and error correction (Didolkar et al., 2024). Drawing on these insights, our Meta-Reasoner framework can be viewed as analogous to dual-process systems: the LLM generates CoT steps akin to System 1, while the Meta-Reasoner provides high-level strategic oversight, analogous to System 2, guiding or redirecting reasoning as needed. This separation of responsibilities balances efficiency with robust problem-solving, allowing the LLM to handle routine inferences and the Meta-Reasoner to intervene for strategic adjustments.

B Experiment Details


Datasets. We mainly focus on tasks that demand complex reasoning and often involve lengthy thinking processes for the correct solutions. These includes (1) 24-point game (Yao et al., 2023); (2) college-level scientific problem from SciBench (Wang et al., 2023); (3) math questions based on theorems from TheoremQA (Chen et al., 2023) and (4) math questions from Math-500 (Lightman et al., 2023) and AIME-2024 (Guan et al., 2025; AI-MO Team, 2024). For SciBench, we focus only on the math-related subsets (i.e., diff, stat, and calc). Detailed explanations for each subset can be found in Wang et al. (2023). For TheoremQA, we only consider the math subset that involves logical reasoning.

Model	Method	Token Usage	Inference Time (s)
qwen-3-8B	Meta-Reasoner	1728.9 ± 42.3	31.70 ± 1.24
qwen-3-8B	MACM	2266.78 ± 58.1	41.35 ± 1.87
qwen-3-8B	ToT (b=5)	2535.72 ± 67.4	46.17 ± 2.15
qwen-3-8B	Best of N	2497.3 ± 63.2	45.48 ± 2.03
qwen-3-8B	Zero-shot	153.68 ± 4.2	3.47 ± 0.18
o1-preview	Zero-shot	3534.64 ± 128.5	44.99 ± 2.67
o1-mini	Zero-shot	2766.24 ± 95.3	17.05 ± 1.12
meta-llama-3.1-8B	Meta-Reasoner	1728.9 ± 38.7	8.93 ± 0.41
meta-llama-3.1-8B	MACM	2266.78 ± 54.2	11.62 ± 0.53
meta-llama-3.1-8B	ToT (b=5)	2535.72 ± 61.8	12.97 ± 0.62
meta-llama-3.1-8B	Best of N	2497.3 ± 59.4	12.78 ± 0.58
meta-llama-3.1-8B	Zero-shot	153.68 ± 3.9	1.06 ± 0.05
gpt-4o-mini	Meta-Reasoner	1728.9 ± 45.1	18.60 ± 0.89
gpt-4o-mini	MACM	2266.78 ± 59.7	24.30 ± 1.15
gpt-4o-mini	ToT (b=5)	2535.72 ± 68.3	27.15 ± 1.34
gpt-4o-mini	Best of N	2497.3 ± 65.8	26.74 ± 1.28
gpt-4o-mini	Zero-shot	153.68 ± 4.5	1.92 ± 0.09
ds-r1-distill-qwen-7B	Meta-Reasoner	1728.9 ± 41.5	17.20 ± 0.82
ds-r1-distill-qwen-7B	MACM	2266.78 ± 56.3	21.26 ± 1.03
ds-r1-distill-qwen-7B	ToT (b=5)	2535.72 ± 65.1	23.29 ± 1.18
ds-r1-distill-qwen-7B	Best of N	2497.3 ± 62.7	23.00 ± 1.12
ds-r1-distill-qwen-7B	Zero-shot	153.68 ± 4.1	5.33 ± 0.28
ds-r1-distill-llama-8B	Meta-Reasoner	1728.9 ± 43.8	31.49 ± 1.52
ds-r1-distill-llama-8B	MACM	2266.78 ± 57.9	40.98 ± 1.96
ds-r1-distill-llama-8B	ToT (b=5)	2535.72 ± 66.7	45.72 ± 2.21
ds-r1-distill-llama-8B	Best of N	2497.3 ± 64.3	45.04 ± 2.15
ds-r1-distill-llama-8B	Zero-shot	153.68 ± 4.3	3.70 ± 0.19

Table 9: Inference Compute Across Different Methods (Mean ± StdDev over 3 runs)

Training Details. We collect the training data for each task: (1) for the 24-point game, we random sample 50 queries from 4nums.com specifically excluding problems in ranks 901-1000 which were reserved for testing; (2) for TheoremQA, we randomly sample 30 mathematical reasoning queries from the dataset; (3) for SciBench, we randomly sample 30 queries from differential subsets including diff, stat, and calc from the entire dataset. These samples were used to iteratively update the LinUCB parameters for both fixed ($K=3$ or $K=5$ strategies) and dynamic strategy settings. We configure the training process using deterministic generation ($n=1$, $Top_k=1$, $temperature=0$) with specific `max_token` limits for CoT generation (512), meta-reasoner feedback (256), progress reports (512), and reward model outputs (4). The LinUCB exploration parameter was set to $c=0.2$. Experiments were ran for a maximum of $T=30$ iterations for the 24-point game, scibench, and $T=100$ for TheoremQA, with MAB parameters updated after each iteration based on a reward function that weighted objective completion (40%), progress quality (30%), efficiency (15%), and strategy alignment (15%). The reward function prompt is detailed in Appendix §F.

Baselines. We consider several established prompting methods as baselines as follows: (1) Chain-of-thought (CoT) (Wei et al., 2022): A prompting technique that encourages models to



(a) Game of 24 Task.

(b) TheoremQA Task

Figure 5: Token Usage on Game-of-24 (Yao et al., 2023) and TheoremQA (Chen et al., 2023) Tasks.

Model	Method	Game-of-24	TheoremQA	Diff (%)	Stat (%)	Calc (%)
gpt-4o-mini	Meta-reasoner	89	84.13	60.32	73.64	80.23
gpt-4o-mini	HiAR-ICL (Wu et al., 2024)	87	83.48	57.42	70.12	77.93
gpt-4o-mini	Evo-Prompt (Guo et al., 2023)	82	81.28	55.28	67.32	76.53
gemini-exp-1206	Meta-reasoner	94	86.32	65.47	79.42	82.77
gemini-exp-1206	HiAR-ICL (Wu et al., 2024)	88	84.41	57.76	75.92	80.23
gemini-exp-1206	Evo-Prompt (Guo et al., 2023)	84	80.32	57.32	70.32	78.42

Table 10: Comparison with HiAR-ICL (Wu et al., 2025) and Evo-Prompt (Guo et al., 2025).

generate intermediate reasoning steps to enhance problem-solving capabilities. (2) Self-Consistent Chain of Thought (SC-CoT) (Wang et al., 2022): An extension of CoT that improves reasoning consistency by generating multiple reasoning chains and selecting the most consistent answer. (3) Multi-Chain Reasoning (MCR) (Yoran et al., 2023): enhances SC-CoT by having another LLM to assess and integrate content among the sampled reasoning chains to generate the final consistent answer. (4) Tree of Thoughts (ToT) (Yao et al., 2023): A method that explores multiple reasoning paths in a tree structure, allowing the model to consider various possibilities before arriving at a conclusion by tree search algorithms. (5) Reflexion (Shinn et al., 2023): A framework that enables models to reflect on their reasoning process, iteratively refining their answers based on feedback. (6) MACM (Lei et al., 2024): A multi-agent system to refine the reasoning based on iterative condition mining. (7) HiAR-ICL (Wu et al., 2025): A high-level automated reasoning paradigm for in-context learning that constructs abstract thinking patterns using Monte Carlo Tree Search (MCTS) on atomic actions, dynamically matching them to problems via a cognitive complexity framework to reduce reliance on specific examples. (8) Evo-Prompt (Guo et al., 2025): A discrete prompt optimization framework that integrates LLMs with evolutionary algorithms, leveraging LLMs to generate coherent prompt variants through operators like crossover and mutation for iterative improvement without gradients.

Backbone Models. We consider both LLMs and the recent Large Reasoning Models (LRMs) for

Parameter Variation	Game-of-24 Acc. (%)	TheoremQA Acc. (%)
<i>Default: $w_1 = 0.5, w_2 = 0.5, \alpha = 0.1, \beta = 0.8$</i>		
Default	89.0 \pm 1.2	84.1 \pm 1.5
<i>Varying w_1 and w_2 ($\alpha = 0.1, \beta = 0.8$)</i>		
$w_1 = 0.3, w_2 = 0.7$ (Adherence emphasis)	87.5 \pm 1.4	83.2 \pm 1.6
$w_1 = 0.7, w_2 = 0.3$ (Correctness emphasis)	88.2 \pm 1.3	85.3 \pm 1.4
$w_1 = 0.9, w_2 = 0.1$ (Heavy correctness)	87.8 \pm 1.5	84.8 \pm 1.7
<i>Varying α ($w_1 = 0.5, w_2 = 0.5, \beta = 0.8$)</i>		
$\alpha = 0.05$ (Low cost penalty)	89.4 \pm 1.1	84.5 \pm 1.4
$\alpha = 0.2$ (Moderate penalty)	88.6 \pm 1.3	83.7 \pm 1.5
$\alpha = 0.5$ (High penalty)	84.2 \pm 1.6	79.8 \pm 1.8
<i>Varying β ($w_1 = 0.5, w_2 = 0.5, \alpha = 0.1$)</i>		
$\beta = 0.6$ (Balanced trade-off)	88.1 \pm 1.2	83.4 \pm 1.5
$\beta = 0.5$ (Efficiency emphasis)	86.7 \pm 1.4	82.0 \pm 1.7
$\beta = 0.9$ (Progress emphasis)	89.3 \pm 1.1	84.6 \pm 1.4

Table 11: Sensitivity analysis results on reward function weights.

our experiments. For the LLMs, we consider the closed-source models like gpt-4o, gpt-4o-mini (between Nov 2025 to Jan 2025) from OpenAI, and open-sourced models like meta-llama-3.1-8B-instruct from Meta, qwen-3-8B from Alibaba, phi-4 from Microsoft, gemini-experimental-1206 from Google, ds-r1-distill-llama-8B, ds-r1-distill-qwen-7B from Deepseek. For the LRMs, we consider the closed-source models like o1, o1-mini (In case we cannot break down the generation of o1 models through APIs, we cannot properly inject our meta-reasoner with o1-series models; we only provide the IO results for references). For the feature extraction mentioned in §3.3, we use text-embedding-3-small from OpenAI as the embedding model. To ensure the reproducibility of the experiments, we set temperature = 0.7 and top_p = 1.0 for all models. We use the API service from OpenAI² and OpenRouter³ for our experiments which host detailed snapshots of the utilized model versions.

Judge Model Analysis. To justify our choice of gemini-2.5-flash as the reward model, we compared it against a reasoning-capable model, gemini-2.5-pro. We randomly sampled 50 cases from Game-of-24 and TheoremQA to test. The results in Table 12 show that the lightweight judge is sufficient for evaluating progress reports, offering the best trade-off between efficiency and performance.

C Sensitivity Analysis of Reward Function Weights

To address concerns about the choice of hyperparameters in the reward function (Section 3.3.1), we perform a sensitivity analysis on the main weights. These include w_1 and w_2 (weights for correctness

Judge Model	Game-24 Acc	TheoremQA Acc	Cost/1K calls
Gemini-2.5-Flash	88.0%	83.5%	\$0.015
Gemini-2.5-Pro	88.5%	84.1%	\$0.125
Difference	+0.5%	+0.6%	+733%

Table 12: Comparison of Judge Models with Reasoning Mode. Using a stronger reasoning model yields negligible accuracy gains (< 1%) but significantly increases cost and latency.

C_c and adherence C_a in solution progress S_p), α (cost coefficient for resource usage R_u), and β (weight balancing S_p and R_u in the total reward R). In the main experiments, we use the following default values: $w_1 = w_2 = 0.5$ (equal emphasis on correctness and adherence), $\alpha = 0.1$ (moderate penalty for extra steps), and $\beta = 0.8$ (strong preference for progress over cost, since accuracy is the primary objective). We select these defaults using a small grid search on a validation subset comprising 10% of the Game-of-24 dataset. The goal is to favor accurate solutions while still encouraging efficient inference.

We then evaluate the robustness of the meta-reasoner under different hyperparameter configurations. We vary the weights and measure accuracy on two benchmarks: Game-of-24 and TheoremQA. For each configuration, we run the full meta-reasoner pipeline (with GPT-4o-mini as the backbone) using 5 random seeds. We report mean accuracy and standard deviation. To isolate the effect of the reward function, we use the fixed contextual bandit variant and disable dynamic strategy generation. We consider three groups of variants:

1. Vary w_1 and w_2 while keeping α and β fixed (to test the balance between correctness and adherence).
2. Vary α (to test sensitivity to penalties on computational cost).

²<https://openai.com/>

³<https://openrouter.ai/>

- Vary β (to test the trade-off between progress and efficiency in the total reward).

Method	TheoremQA		Game-of-24	
	Acc (%)	Time (s)	Acc (%)	Time (s)
Full Method	84.13	18.60	89.0	43.8
w/o MAB (Direct Selection)	80.74	18.17	82.0	42.9
Difference	-3.39	-0.43s (-2.3%)	-7.0	-0.9s (-2.1%)

Table 13: Impact of CMAB on computational overhead and performance using gpt-4o-mini. The MAB component introduces negligible latency (<3%) while significantly boosting accuracy.

Table 11 reports the results. Performance remains stable across a broad range of weight values. Even in extreme settings, accuracy degrades by at most 5-7% relative to the default configuration. For example, increasing the emphasis on correctness ($w_1 = 0.7$, $w_2 = 0.3$) yields a small accuracy gain on TheoremQA (+1.2%), which likely reflects that task’s focus on logical precision, while equal weighting works best on Game-of-24. A large cost coefficient (e.g., $\alpha = 0.5$) reduces the number of reasoning steps by 20-25% but harms accuracy on long-horizon tasks such as TheoremQA, which supports the use of moderate efficiency penalties. A lower value $\beta = 0.5$ shifts the reward towards efficiency, giving similar accuracy with 15-20% fewer tokens on average; this configuration can be useful in resource-constrained settings. Overall, these findings show that the benefits of the meta-reasoner do not rely on precise hyperparameter tuning. The main gains come from the architecture itself, in particular dynamic strategy selection via CMABs and the progress-report interface. Even with sub-optimal weights, our method outperforms strong baselines such as MACM (80 on Game-of-24) and Reflexion (74.32% on TheoremQA).

D Stability of Dynamic Strategy Generation.

Dynamic expansion and refinement of the strategy set (Section 3.3) increase adaptability. At the same time, they introduce a risk of instability, since naive strategy generation can be random or ungrounded. We address this risk with three stabilizing mechanisms.

(1) *Initial stable foundation.* We begin with a curated set of verified strategies (e.g., Table 1). Examples include $g_1 =$ “Pause to clarify and disambiguate reasoning” and $g_2 =$ “Decompose the task into sub-tasks.” This set provides a strong baseline. During inference, the LLM proposes new strategies

g_t based on the current progress report P_t . The proposal is constrained by the task context and the report content, so the system does not add arbitrary or irrelevant strategies.

(2) *Exploration–exploitation balance.* The CMAB uses an ϵ -greedy policy to balance exploration and exploitation. At meta-reasoning round t , the probability of exploring a newly added arm $a_t \in G_t \setminus G_{t-1}$ is $\epsilon_t = \frac{1}{t}$. Thus, exploration is frequent at early rounds and gradually decays. We filter suboptimal strategies during inference based on their empirical reward $R_t(a_t, P_t)$. On the Game-of-24 task, this dynamic bandit achieves 89% accuracy, whereas fixed strategy sets yield 65%–72% accuracy (Table 7). These results indicate that the policy effectively prioritizes viable strategies.

(3) *Reward-driven feedback.* The reward function $R : G_t \times P_t \rightarrow [0, 1]$ (defined in Section 3.3) provides online feedback on strategy quality. Strategies that consistently yield low rewards (e.g., $R_t < 0.3$) are rapidly down-weighted and effectively removed from future consideration. Figure 3 shows that the cumulative reward $\sum_{i=1}^t R_i$ increases over time as the controller incorporates new, useful strategies. This pattern reflects a stable and improving reasoning process.

These three mechanisms act together. Context-constrained generation anchors new strategies in the current reasoning state. Bandit optimization filters and reorders strategies based on observed rewards. Reward-driven feedback then refines the strategy set over time. Empirical results on Game-of-24 and TheoremQA (Table 7) show that dynamic strategy generation yields 17%–24% gains over static strategies, while maintaining stable behavior.

E Computational Overhead of CMAB

To assess the latency introduced by the meta-reasoning module, we compare the computational cost of the CMAB with the cost of base LLM inference. The CMAB uses the LinUCB algorithm and text-embedding-3-small for feature extraction. Both components run in a few milliseconds per meta-reasoning step. In contrast, the main cost comes from the LLM, which generates hundreds of tokens and often takes several seconds. The additional latency from CMAB is therefore small relative to the total inference time. Table 13 reports a trade-off analysis with gpt-4o-mini. Adding the CMAB increases total inference time by less than 3% (about 0.4–0.9 seconds for tasks that take

18–44 seconds). At the same time, it improves accuracy by 3–7 percentage points. These results show that the meta-reasoner using CMAB incurs minimal computational overhead while providing clear performance gains.

F Prompt List

Prompt for Progress Report (§3.2)

You are an advanced AI summarizer with expertise in extracting and condensing key insights from recent developments. Your goal is to create a concise progress report based on the provided information.

Read the task description and the chain of thoughts generated so far. Please ignore the <examples> section which is only for the demonstration of the task.

<progress>
SOLUTION
</progress>

And then complete the following template:

Current Attempts:

[Insert the list of previous attempts here]

Analysis Instructions:

1. Systematically review each attempted step
2. Identify patterns in the current solution attempts
3. Provide observations regarding:
 - Recurring strategies
 - Missed opportunities
 - Potential promising approaches
 - Any mathematical observations about the number combination

Output Format:

- Provide a structured analysis
- Include bullet points for key observations

Constraints:

- Use clear, logical reasoning
- Focus on mathematical problem-solving approaches
- Avoid random guessing
- Make the analysis short and to the point (around 6-7 sentences)

Prompt for Meta-Reasoner (Dynamic Bandit Generation (§3.3.3))

You are a Meta-reasoner, tasked with analyzing the reasoning process of another agent and providing guidance for its further steps. Your goal is to improve the efficiency and effectiveness of that agent's problem-solving approach.

Review the task description and the summary of the recent reasoning progress below:
PROGRESS_REPORT

Provide feedback in the following format:

- Reflection: What is the current strategy of the agent to solve the task? Has the agent made sufficient progress? Are there any mistakes or misconceptions in the intermediate steps? Is the agent taking unnecessary detours or repeating steps?
- Fact Check: Are the agent's statements accurate and relevant to the task? Are there any logical errors or incorrect assumptions?
- Thought: What are the key insights or strategies that the agent should focus on? Are there alternative methods or perspectives that could be beneficial?
- Action: The action to take

Make your response precise and focused without unnecessary details.

Prompt for CoT Generation (§3.3)

You are an AI assistant tasked with generating steps to solve mathematical problems. Your role is to read a task description, consider the current step (if any), and generate the next logical step towards solving the problem. You will also receive feedback from a Meta-reasoner, which you should take into account when determining your next step.

Here is the task description:

```
<task_description>
TASK_DESCRIPTION
</task_description>
```

The process will work as follows:

1. You will be given the current step (if any) in the problem-solving process.
2. You will also receive feedback from the Meta-reasoner about the previous step.
3. Your job is to generate the next logical step towards solving the problem, taking into account the task description, the current step, and the Meta-reasoner's feedback.

To generate the next step:

1. Carefully analyze the task description, the current step (if any), and the Meta-reasoner's feedback.
2. If the Meta-reasoner suggests backtracking, consider how to modify or correct the previous step.
3. If the Meta-reasoner suggests continuing, think about the logical progression from the current step.
4. If the Meta-reasoner suggests changing strategy, brainstorm alternative approaches to the problem.
5. Formulate a clear, concise next step that moves towards solving the problem.

Your response should be a single, well-thought-out step that progresses the problem-solving process. Do not solve the entire problem at once; focus on generating just the next logical step.

Please provide your next step within `<next_step>` tags. Before giving your next step, explain your reasoning within `<reasoning>` tags. Explicitly state whether the problem is solved or not before providing the next step or final answer.

If you believe there has been enough progress to solve the problem completely, generate the final answer in the form of `boxedanswer` at the end of your response. The answer should be a numerical value.

Your response should follow this structure:

```
<reasoning>
[Explain your thought process here, considering the task description, current step, and
Meta-reasoner feedback (make sure to address any issues raised by the Meta-reasoner).
The reasoning should be clear, logical, and directly related to the problem-solving process.]
</reasoning>
```

```
<next_step>
[Provide the next logical step here]
</next_step>
```

[State whether the problem is solved or not]

[If the problem is solved] Return only the Final answer:
`boxednumerical_value`

Remember to focus on generating just the next logical step, not solving the entire problem at once (unless you've reached the final solution). Your explanation and step should be clear, concise, and directly contribute to solving the mathematical problem at hand.

Here is the current step (if this is the first step, this will be empty):

```
<current_step>
CURRENT_STEP
</current_step>
```

And here is the feedback from the Meta-reasoner (if this is the first step, this will be empty):

```
<meta_reasoner_feedback>
META_REASONER_FEEDBACK
</meta_reasoner_feedback>
```

Prompt for Progress Evaluation (§3.2)

You are an impartial evaluator tasked with assessing the progress of a reasoning process toward solving a given task objective. Your evaluation must be based strictly on the provided reward function components. Do not favor any particular output or introduce bias—evaluate objectively.

Inputs:

Task Objective (G_t): [INSERT TASK OBJECTIVE HERE] // e.g., the original user query or problem statement

Current Progress (P_t): [INSERT CURRENT REASONING/PROGRESS HERE] // e.g., the model's accumulated reasoning steps, partial solution, or plan up to this point

Number of Reasoning Steps (N_s): [INSERT NUMBER OF STEPS HERE] // e.g., the count of iterative reasoning steps taken so far

Weights and Coefficients:

- w₁ (weight for correctness): 0.5
- w₂ (weight for adherence): 0.5
- alpha (cost coefficient for resource usage): 0.1
- beta (trade-off balance): 0.8

Evaluation Criteria:

1. Correctness (C_c): Score on a scale of 0.0 to 1.0 how accurate and logically sound the current progress is toward fully solving the task objective. Consider factual accuracy, logical consistency, and advancement toward a complete solution. 0.0 means no progress or entirely incorrect; 1.0 means perfectly correct and on track for completion.
2. Adherence (C_a): Score on a scale of 0.0 to 1.0 how well the current progress follows the task objective's constraints, requirements, and guidelines (e.g., format, scope, ethical considerations). 0.0 means complete disregard; 1.0 means full compliance.
3. Solution Progress (S_p): Compute as $S_p = (w_1 * C_c) + (w_2 * C_a)$.
4. Resource Usage (R_u): Compute as $R_u = -\alpha * N_s$. This penalizes excessive steps for efficiency.
5. Total Reward (R): Compute as $R = (\beta * S_p) + ((1 - \beta) * R_u)$.

Output Format:

Respond only in the following strict JSON structure. Do not include any additional text, explanations, or commentary outside this JSON.

```
{
  "C_c": <float, your score for correctness>,
  "C_a": <float, your score for adherence>,
  "S_p": <float, computed solution progress>,
  "R_u": <float, computed resource usage>,
  "R": <float, total reward>,
  "brief_rationale": "<A concise 1-2 sentence explanation for C_c and C_a scores only.>"
}
```