

ROUTE-AND-REASON: SCALING LARGE LANGUAGE MODEL REASONING WITH REINFORCED MODEL ROUTER

006 **Anonymous authors**

007 Paper under double-blind review

ABSTRACT

013 Chain-of-thought has been proven essential for enhancing the complex reasoning
 014 abilities of Large Language Models (LLMs), but it also leads to high computational
 015 costs. Recent advances have explored the method to route queries among multiple
 016 models and proved it as a promising approach. However, previous works directly
 017 operate at the task level, i.e., assigning user queries to suitable LLMs, which does
 018 not allow hybrid LLMs to truly collaborate on finer-grained sub-tasks. Collabora-
 019 tion at the level of intermediate reasoning steps (thoughts) could enable more
 020 efficient coordination, but it also poses significant challenges for router scheduling,
 021 placing immense demands on the quality of task decomposition and the precision of
 022 the router. To address this, we propose **R2-Reasoner**, a novel framework centered
 023 around **a Reinforced Model Router** designed to efficiently scale LLM reasoning.
 024 This router orchestrates collaboration across 9 heterogeneous models, of whom the
 025 parameter scale ranges from less than 1B to hundreds of billions, by first breaking
 026 down a complex query into subtasks with a decomposer, and then assigning each
 027 subtask to the optimal model with a subtask allocator, balancing performance with
 028 cost. To train this router involves a two-stage alternating process for the decomposer
 029 and the allocator, integrating supervised fine-tuning with reinforcement learning
 030 to enable effective self-supervised refinement. Extensive experiments across six
 031 challenging reasoning benchmarks demonstrate that R2-Reasoner reduces API
 032 costs by 84.46% compared with state-of-the-art baselines while maintaining com-
 033 petitive reasoning accuracy. Our framework paves the way for the development
 034 of more scalable and efficient reasoning systems. Our code is open-source at
 035 https://anonymous.4open.science/r/R2_Reasoner.

1 INTRODUCTION

037 Chain-of-Thought (CoT, (Wei et al., 2022)) reasoning has endowed large language models (LLMs)
 038 with significantly enhanced reasoning capabilities. The reasoning ability of LLMs has evolved from
 039 prompting-based sequential thoughts to reinforcement learning–driven long-chain reasoning (OpenAI,
 040 2024; Guo et al., 2025; Lightman et al., 2023; Snell et al., 2024; Wang et al., 2023; Chen et al., 2024b),
 041 developing into *test-time scaling* as a paradigm, though with significant computational cost (Wu et al.,
 042 2024). To mitigate the vast increase in computational overhead, **model router** has been introduced
 043 to route queries across models according to problem difficulty, model capability and associated
 044 cost. This strategy is recognized as an effective means of balancing the enhancement of reasoning
 045 performance with cost control. Its recent deployment in GPT-5 (OpenAI, 2025) further demonstrates
 046 the great potential of this approach.

047 Recent studies have increasingly explored model routers in various scenarios. One line of research
 048 aims to select one or more models that are most suitable for each task from a knowledge coverage
 049 perspective (Feng et al., 2024; Zhang et al., 2025; Dekoninck et al., 2024; Chen et al., 2024c).
 050 This approach can be viewed as a form of LLM ensembling, motivated by the observation that
 051 different LLMs exhibit complementary strengths in different knowledge domains. While effective for
 052 knowledge-intensive tasks such as factual QA, these methods are limited when applied to multi-step
 053 complex reasoning (e.g., mathematical derivations), and they seldom explicitly optimize for inference
 cost efficiency. Another line of work focuses on device-cloud collaboration, where local lightweight

054 small language models (SLMs) and cloud-based LLMs are coordinated such that simpler tasks are
 055 routed to SLMs, while more complex tasks are escalated to LLMs (Chen et al., 2024a; Shao et al.,
 056 2025a; Li et al., 2019; Hao et al., 2024). However, operating at the task level often results in overly
 057 coarse routing granularity, making accurate routing decisions challenging and introducing additional
 058 overhead.

059 To address these limitations, we revisit the problem of model routing from the perspective of sub-
 060 tasks. Even complex reasoning problems often comprise relatively simple sub-tasks, which can be
 061 effectively resolved by more computationally efficient small-scale language models (SLMs). If these
 062 simpler “thoughts” can be accurately identified and delegated to such SLMs, while reserving the more
 063 complex, capability-intensive sub-problems for larger LLMs, the overall cost can be substantially
 064 reduced. This hierarchical approach to model utilization aligns naturally with typical data center
 065 deployment scenarios, where a diverse set of models with varying capabilities is often available,
 066 enabling dynamic allocation based on sub-task requirements.

067 Nevertheless, implementing such a framework faces two core challenges. First, high-quality task
 068 decomposition, splitting the overall problem into coherent, solvable sub-tasks, is non-trivial (Wies
 069 et al., 2023; Zhou et al., 2022), as poor decomposition can produce erroneous intermediate steps or
 070 inefficient work allocation, undermining both outcomes and efficiency (Zhu et al., 2023; Zheng et al.,
 071 2023). Second, determining the difficulty of each sub-task is challenging but critical for assigning the
 072 right model; errors may overload smaller models or waste larger ones, reducing inference efficiency,
 073 accuracy, and the potential cost savings of this approach.

074 To overcome these challenges, we propose **R2-Reasoner**, a framework that leverages a *Reinforced*
 075 *Model Router* to efficiently scale LLM reasoning. As the core component, the Router operationalizes
 076 task decomposition and subtask allocation as two distinct yet interconnected LLMs: the Task
 077 Decomposer generates a structured sequence of sub-tasks from a complex input, while the Subtask
 078 Allocator assigns each subtask to the most suitable model, ranging from lightweight SLMs to powerful
 079 LLMs, based on estimated difficulty. By explicitly separating decomposition and allocation, R2-
 080 Reasoner enables fine-grained, scalable collaboration across heterogeneous models, optimizing both
 081 reasoning accuracy and computational efficiency.

082 To fully unlock the potential of the *Model Router*, we develop a staged reinforcement learning
 083 pipeline that progressively refines its decision-making capability. We decouple the joint training of
 084 the Decomposer and Allocator, two core LLMs, into an alternating iterative process, avoiding the non-
 085 differentiability and gradient blockage in end-to-end updates across multiple LLMs. This approach
 086 combines supervised fine-tuning on task-specific data with Group Relative Policy Optimization
 087 (GRPO) in a multi-stage pipeline, enabling stable and coordinated policy improvement through
 088 self-supervised feedback. The framework requires no additional human annotation and enhances
 089 adaptability in dynamic real-world scenarios.

090 Extensive evaluations across 6 benchmarks validate the efficacy of our framework. The results
 091 demonstrate a substantial reduction in inference costs, achieving an 84.46% decrease in API expenses
 092 while maintaining reasoning performance competitive with strong baseline methods and even im-
 093 proving average accuracy by 3.73%. Additionally, we conduct further experiments to demonstrate
 094 that R2-Reasoner exhibits strong generalization, capable of directly adapting to previously unseen
 095 models. Moreover, our framework supports a flexible and controllable trade-off between accuracy
 096 and inference cost, enabling practical deployment across diverse budget scenarios. In summary, our
 097 key contributions are:

- 098 • We propose **R2-Reasoner**, a novel framework centered around a *Reinforced Model Router*
 099 designed to efficiently scale LLM reasoning at test-time. This framework facilitates fine-
 100 grained, collaborative reasoning by decomposing complex tasks and strategically allocating
 101 subtasks across a diverse pool of heterogeneous models.
- 102 • We introduce a staged training pipeline to optimize the *Model Router*. This iterative training
 103 strategy not only enables the router to iteratively refine its performance but also circumvents
 104 the non-differentiability that arises in end-to-end gradient propagation between two LLMs.
- 105 • We conduct extensive experiments on six complex reasoning benchmarks, demonstrating
 106 that R2-Reasoner can substantially reduce reasoning costs while maintaining high reasoning
 107 accuracy, thereby paving the way for more scalable test-time scaling.

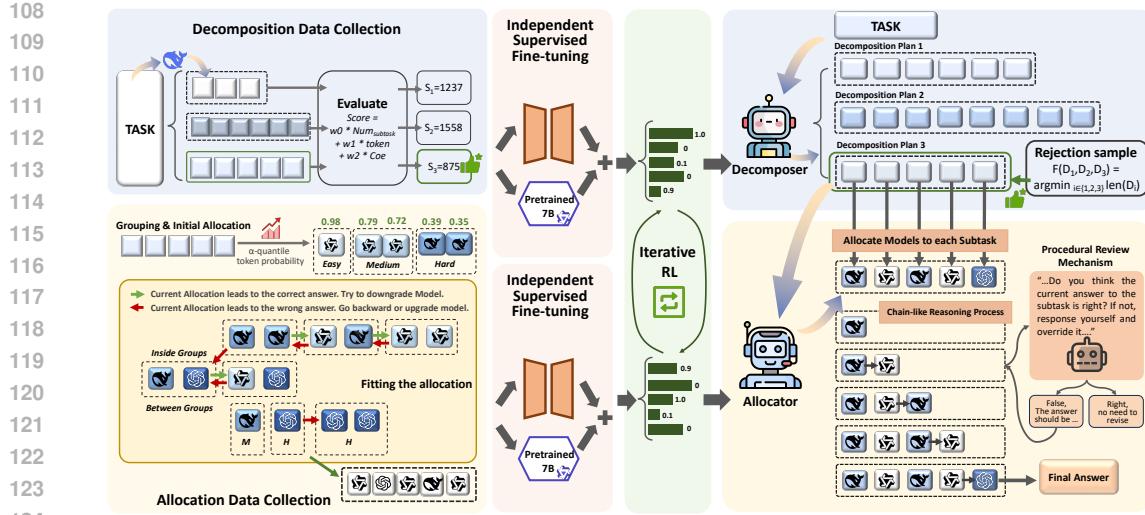


Figure 1: Overview of Our R2-Reasoner Framework

2 RELATED WORKS

2.1 TASK DECOMPOSITION AND MULTI-STEP REASONING

The chain-of-thought (CoT) prompting technique (Wei et al., 2022) has emerged as a key method for enhancing LLM reasoning, enabling step-by-step inference without additional training. Building on this idea, more advanced paradigms such as tree-of-thought (ToT) (Yao et al., 2023) generalize reasoning into structured sequences of intermediate ‘thoughts.’ Leveraging this notion, task decomposition methods and process reward models (Lightman et al., 2023) have been proposed to guide or supervise individual reasoning steps. Together, these approaches illustrate an emerging paradigm that scales reasoning through both structural decomposition and increased compute (Snell et al., 2024).

2.2 COLLABORATIVE REASONING AMONG LLMs

Recent research has explored several strategies for enabling collaborative reasoning among multiple language models, each with distinct trade-offs. Model partitioning (Li et al., 2019; Cai et al., 2024; Zhang et al., 2024) distributes a single LLM across nodes, but suffers from high communication overhead and limited robustness. Simple referral (Chen et al., 2024a) routes easy queries to small models and harder ones to stronger LLMs, though performance depends on accurately assessing query difficulty. Token correction (Hao et al., 2024) lets an SLM draft outputs while an LLM revises suboptimal tokens, improving quality but incurring extra decoding costs. Despite these advances, existing methods remain constrained by coordination efficiency, accuracy, and scalability, underscoring the need for more adaptive collaboration frameworks.

3 PRELIMINARIES

Problem Definition

We consider a highly general scenario of a large-scale LLM platform, where numerous models are deployed locally on the platform, and some are hosted in the cloud. The goal of the platform is to make comprehensive use of these LLMs to provide users with high-quality inference services at the lowest possible cost. Denote the local deployed SLMs as $\mathcal{M}_{\mathcal{E}} = \{\mathcal{M}_{\mathcal{E}_1}, \mathcal{M}_{\mathcal{E}_2}, \dots, \mathcal{M}_{\mathcal{E}_n}\}$, and the cloud-based LLMs as $\mathcal{M}_{\mathcal{C}} = \{\mathcal{M}_{\mathcal{C}_1}, \mathcal{M}_{\mathcal{C}_2}, \dots, \mathcal{M}_{\mathcal{C}_n}\}$. The user’s original query is restricted to the edge model for task decomposition and allocation, while the resulting sub-tasks can be resolved by either $\mathcal{M}_{\mathcal{E}}$ or $\mathcal{M}_{\mathcal{C}}$. The entire set of reasoning tasks is represented as $\mathcal{T} = \{T_1, T_2, \dots, T_n\}$. Let the reasoning accuracy over the entire task set be denoted as Acc , with the API cost represented by C_{Api} .

For each task T , denote the decomposition process as: $T \rightarrow \{t^1, t^2, \dots, t^k\}$. Based on the decomposed subtasks t^i , the model allocation scheme can be denoted as: $M : t^i \mapsto \{\mathcal{M}_{\mathcal{E}}, \mathcal{M}_{\mathcal{C}}\}$, which prioritizes assigning simple subtasks to on-device SLMs, while invoking the cloud-based LLM for

162 handling complex subtasks. The goal of our optimization is to minimize the discrepancy between the
 163 model’s allocation scheme M and the optimal scheme M^* : $\min |M - M^*|$. The optimal scheme M^*
 164 is derived through a search strategy that maximizes SLM usage while maintaining accuracy. During
 165 the optimization process, as the allocation scheme gradually approaches the optimal solution, the
 166 API cost C_{Api} decreases, while Acc remains well-maintained.
 167

168 4 METHODOLOGY

170 The R2-Reasoner framework is centered around a **Model Router**, which consists of two primary
 171 modules: a **Task Decomposer** ($\mathcal{M}_{\text{decomp}}$) and a **Subtask Allocator** ($\mathcal{M}_{\text{alloc}}$). The Task Decomposer is
 172 engineered to break down complex input tasks T into more manageable, well-structured, and logically
 173 ordered subtasks $\{t^1, t^2, \dots, t^k\}$. Following this, the Subtask Allocator strategically assigns each
 174 subtask t^i to the most suitable model from a heterogeneous pool ($\mathcal{M}_{\text{pool}} = \mathcal{M}_{\mathcal{E}} \cup \mathcal{M}_{\mathcal{C}}$, comprising
 175 models of diverse capabilities). This allocation process is driven by the estimated difficulty of each
 176 subtask, aiming to strike an optimal balance between reasoning fidelity and computational resource
 177 expenditure. The design and training of these interconnected components are detailed below.
 178

179 4.1 GENERATING COHERENT SUBTASK SEQUENCES VIA TASK DECOMPOSER

180 The **Task Decomposer** ($\mathcal{M}_{\text{decomp}}$) serves as the first stage of the Model Router, responsible for
 181 transforming a complex task T into a sequence of logically connected subtasks $\{t^1, t^2, \dots, t^k\}$.
 182 The quality of this decomposition is crucial: redundant or incoherent breakdowns can cause error
 183 propagation, while clear and concise subtasks provide a strong foundation for subsequent allocation.
 184

185 To supervise training, we construct a decomposition dataset $\mathcal{D}_{\text{decomp}}$ using a rejection sampling strat-
 186 egy. For each task, multiple candidate decompositions are generated and then evaluated along three
 187 dimensions: **Conciseness**, assessed by the number of subtasks to avoid both excessive fragmentation
 188 and overly coarse splits. **Practicality**, estimated by the total token cost of solving all subtasks with a
 189 baseline model. **Coherence**, measuring the logical continuity between adjacent subtasks, with fewer
 190 breaks indicating higher quality.
 191

192 These criteria are linearly combined into a weighted score, where lower values correspond to higher-
 193 quality decompositions. A binary correctness signal $C(d) \in \{0, 1\}$ is further incorporated to ensure
 194 that the selected decomposition can solve the original task. When possible, only candidates with
 195 $C(d) = 1$ are retained, and among them the one with the best score is chosen. This guarantees that
 196 $\mathcal{D}_{\text{decomp}}$ contains decompositions that are concise, coherent, and practical while remaining effective
 197 for solving the task. The resulting pairs (T, d^*) are then used to fine-tune $\mathcal{M}_{\text{decomp}}$. More details and
 198 formulas are provided in the Appendix B.1.
 199

200 4.2 STRATEGIC MODEL ASSIGNMENT FOR COLLABORATION VIA SUBTASK ALLOCATOR

201 Once $\mathcal{M}_{\text{decomp}}$ produces a subtask sequence, the **Subtask Allocator** ($\mathcal{M}_{\text{alloc}}$) determines how to
 202 distribute these subtasks across the heterogeneous model pool $\mathcal{M}_{\text{pool}}$. Formally, for each subtask
 203 t^i , it selects a model $M_j \in \mathcal{M}_{\text{pool}}$, yielding an assignment $M_A : t^i \mapsto M_j$. To enable $\mathcal{M}_{\text{alloc}}$ to
 204 learn efficient assignment policies, we construct a high-quality dataset $\mathcal{D}_{\text{alloc}}$ of model allocation
 205 schemes. Rather than relying on hand-crafted heuristics, we employ a systematic search procedure
 206 over the vast space of possible assignments, seeking schemes that minimize resource consumption
 207 while maintaining perfect accuracy. The resulting allocation pairs $(\{t^i\}, M_A^*)$ serve as supervision
 208 signals for training $\mathcal{M}_{\text{alloc}}$ to imitate these cost-effective strategies.
 209

210 However, exhaustive search over all allocations would be prohibitively expensive in both time and
 211 cost. We therefore design a **Grouped Search Strategy** to approximate optimal assignments efficiently.
 212 The process begins by estimating the difficulty of each subtask t^i using the predictive confidence of a
 213 baseline model: if the maximum token probability exceeds a threshold τ_{easy} , the subtask is labeled
 214 as *easy*; if it falls below τ_{hard} , it is labeled as *hard*; otherwise, it is labeled as *medium*. In parallel,
 215 the model pool $\mathcal{M}_{\text{pool}}$ is partitioned into three capability groups: small language models (SLMs),
 medium language models (MLMs), and large language models (LLMs). Each difficulty level is paired
 with the corresponding capability group (*easy*→SLM, *medium*→MLM, *hard*→LLM).

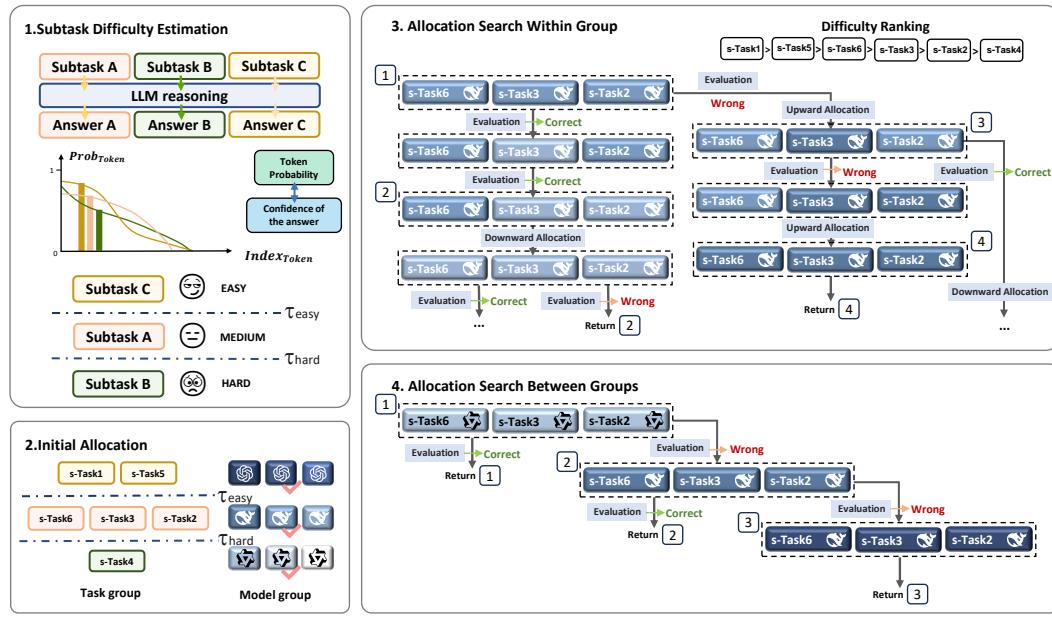


Figure 2: Overview of Our Grouped Search Strategy for Optimal Allocation Scheme

Based on this categorization, an **initial allocation** $M_A^{(0)}$ is obtained by assigning each subtask to the medium-capacity model within its corresponding group. This serves as the starting point for iterative refinement: if the current allocation already achieves correctness ($Acc = 1$), the allocator attempts to replace some models with cheaper ones to reduce cost; if correctness fails, subtasks are escalated to stronger models within the same group, and only if necessary, across groups. The search is bounded by a maximum number of iterations ($N_{\text{iter_alloc}} \leq 20$), after which the resulting allocation M_A^* is accepted. The collection of such $(\{t^i\}, M_A^*)$ pairs constitutes $\mathcal{D}_{\text{alloc}}$, which is then used to train $\mathcal{M}_{\text{alloc}}$. Details of the search algorithm is shown in Algorithm 1.

This strategy enables $\mathcal{M}_{\text{alloc}}$ to learn fine-grained, capability-aware assignment policies that balance accuracy and efficiency. The detailed formulation of the grouped search procedure is deferred to Appendix B.2.

4.3 DUAL-MODULE CO-TRAINING VIA ITERATIVE REINFORCEMENT LEARNING

After the initial SFT of $\mathcal{M}_{\text{decomp}}(\theta_{\text{decomp}})$ and $\mathcal{M}_{\text{alloc}}(\theta_{\text{alloc}})$, We employ a staged RL pipeline to further refine their capabilities and promote synergistic collaboration within the Model Router. In each iteration, one module’s parameters are updated while the other remains fixed, allowing targeted improvements based on task success feedback, which also circumvents the non-differentiability and discontinuities arising from cascading two LLMs, thereby stabilizing training. The primary reward signal is a binary indicator based on the final correctness of the task T :

$$R_{\text{final}}(T, \{t^i\}, M_A) = \begin{cases} 1 & \text{if final answer is correct} \\ 0 & \text{if final answer is incorrect} \end{cases} \quad (1)$$

We adopt Group Relative Policy Optimization (GRPO) as the optimization algorithm for this co-training phase. Training proceeds iteratively for each module:

1. **Updating $\mathcal{M}_{\text{decomp}}(\theta_{\text{decomp}})$:** The decomposer acts as the policy, generating sequences of subtasks $\{t^i\}$ for an input task T . The fixed allocator $\mathcal{M}_{\text{alloc}}(\bar{\theta}_{\text{alloc}})$ assigns models to these subtasks, and the final outcome is used to compute R_{final} . The reward is propagated back to estimate the advantage $\hat{A}_{i,k}$ for decomposition decisions.
2. **Updating $\mathcal{M}_{\text{alloc}}(\theta_{\text{alloc}})$:** The allocator acts as the policy, generating assignments $M_A(t^k)$ for each subtask t^k provided by the fixed decomposer $\mathcal{M}_{\text{decomp}}(\bar{\theta}_{\text{decomp}})$. The final correctness again determines R_{final} , which guides the advantage estimates $\hat{A}_{i,k}$ for allocation choices.

270 This alternating optimization encourages the two modules to progressively adapt to each other,
 271 leading to improved overall reasoning performance. The detailed algorithmic design can be found in
 272 Appendix B.3.

274 4.4 END-TO-END REASONING WORKFLOW AT TEST TIME

275 With the R2-Reasoner’s Task Decomposer ($\mathcal{M}_{\text{decomp}}$) and Subtask Allocator ($\mathcal{M}_{\text{alloc}}$) trained through
 276 SFT and the iterative RL pipeline, the framework can be deployed for inference. For a user query
 277 Q_{user} , the workflow is as follows: (1) **Task Decomposition**: The query Q_{user} is first processed by
 278 the fine-tuned Task Decomposer: $\{t^1, \dots, t^k\} = \mathcal{M}_{\text{decomp}}(Q_{\text{user}})$. (2) **Subtask Allocation**: The
 279 resulting sequence of subtasks $\{t^1, \dots, t^k\}$ is then passed to the fine-tuned Subtask Allocator for
 280 strategic model assignment: $M_A = \mathcal{M}_{\text{alloc}}(\{t^1, \dots, t^k\})$, where $M_A(t^i) \in \mathcal{M}_{\text{pool}}$ is the model
 281 assigned to subtask t^i . (3) **Subtask Execution**: Each subtask t^i is executed sequentially by its
 282 assigned model $M_A(t^i)$. The output of subtask t^i can serve as input to the subsequent subtask
 283 t^{i+1} . (4) **Result Integration**: The results from the executed subtasks are sequentially integrated to
 284 formulate the final answer A_{final} .

285 To flexibly adapt to scenarios with different cost budgets, achieve a controllable accuracy–cost trade-
 286 off, and enhance reasoning robustness, we introduce an optional **Procedural Review Mechanism**
 287 (PRM). Let $\mathcal{M}_{\text{strong}}$ denote a high-capability model (e.g., a frontier LLM from $\mathcal{M}_{\text{pool}}$) and $\mathcal{M}_{\text{thresh}}$
 288 a pre-defined threshold model representing a minimum capability level. For each subtask t^j , let r_j
 289 be the output generated by its initially assigned model $M_A(t^j)$. If $M_A(t^j)$ is below the threshold
 290 $\mathcal{M}_{\text{thresh}}$, the output will be verified and potentially refined: $r_j^{\text{final}} = \text{PRM_Verify}(\mathcal{M}_{\text{strong}}, r_j)$. The
 291 PRM_Verify function utilizes $\mathcal{M}_{\text{strong}}$ to assess the correctness of r_j . If r_j is deemed incorrect or
 292 suboptimal, $\mathcal{M}_{\text{strong}}$ provides a corrected or refined response r'_j ; otherwise, $r_j^{\text{final}} = r_j$. This r_j^{final}
 293 is then used for all subsequent reasoning steps. This mechanism allows targeted quality control,
 294 preserving accuracy while maintaining the cost-efficiency of allocation.

297 5 EXPERIMENTS

300 5.1 EXPERIMENTAL SETUP

301 **Benchmarks**

302 We evaluate our framework on six widely-used open-source benchmarks: (1) **P3** (Schuster et al.,
 303 2021) for program synthesis, (2) **SCAN** (Lake & Baroni, 2018) for language-driven navigation,
 304 (3) **MATH** (Hendrycks et al., 2021) and **CHAMP** (Mao et al., 2024) for solving challenging math
 305 problems, and (4) **CSQA** (Talmor et al., 2018) and **MuSiQue** (Trivedi et al., 2022) for commonsense
 306 reasoning. For each benchmark, we manually annotate a small set of samples for in-context learning
 307 in task decomposition and select another 200 tasks as the test set. Detailed descriptions and dataset
 308 statistics are provided in the Appendix C.2.

309 **Baselines** Considering the scenario of collaborative reasoning, we establish six baselines. (1) **CoT** (Wei et al., 2022): The CoT (Chain of Thought) method asks a single LLM to solve a task by
 310 decomposing the original task into a sequence of sub-tasks and answering these sub-tasks sequentially.
 311 (2) **ToT** (Yao et al., 2023): The ToT (Tree of Thoughts) method, based on the framework of CoT,
 312 prompts multiple answers ($N = 2$) for each sub-task, and retain the best answer by utilizing a scoring
 313 method. It also only deploys one certian LLM. (3) **DataShunt** (Chen et al., 2024a): The Datashunt
 314 method dynamically selects between a SLM and a LLM to finish the task. The method first evaluates
 315 the difficulty of the given task, and allocate the task to either SLM or LLM to solve utilizing the
 316 CoT method. (4) **AutoMix** (Aggarwal et al., 2025): The AutoMix method consists of a few-shot
 317 self-verification mechanism conducted by SLM to evaluate the confidence toward an answer from
 318 SLM and a router that strategically routes queries to LLM based on the confidence. (5) **DoT** (Shao
 319 et al., 2025b): The DoT method decomposes a task into subtasks, builds a dependency graph, and
 320 allocates subtasks to SLMs or LLMs using a Plug-and-Play Adapter on SLMs. This framework
 321 enables efficient edge-cloud collaborative reasoning. (6) **Router-R1** (Zhang et al., 2025): The
 322 Router-R1 method chooses an language model as the router itself, interweaving thinking process by
 323 the router with routing process by the routed models, and integrates every response into the context.

Model	Program Synthesis		Language-Driven Navigation		Math Problem Solving				Commonsense Reasoning			
	P3		SCAN		MATH		CHAMP		CSQA		MuSiQue	
	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}
COT (GPT-4o)	42%	<u>4.45¢</u>	<u>68%</u>	<u>2.75¢</u>	51.5%	<u>5.34¢</u>	55.5%	<u>4.45¢</u>	80%	<u>3.60¢</u>	57%	<u>0.85¢</u>
TOT (GPT-4o)	38%	<u>14.55¢</u>	52%	<u>9.82¢</u>	63%	<u>9.97¢</u>	57%	<u>11.65¢</u>	82%	<u>20.50¢</u>	59%	<u>2.45¢</u>
COT (Llama 3-8B)	5.5%	-	17%	-	10%	-	19%	-	70%	-	38%	-
TOT (Llama 3-8B)	5.5%	-	13%	-	29.5%	-	25%	-	68.5%	-	31%	-
DataShunt	14%	<u>2.45¢</u>	23.5%	<u>1.72¢</u>	16%	<u>1.66¢</u>	34%	<u>2.98¢</u>	73%	<u>1.28¢</u>	47%	<u>0.46¢</u>
AutoMix	14%	<u>0.04¢</u>	43%	<u>0.12¢</u>	44%	<u>0.03¢</u>	44%	<u>0.34¢</u>	66%	<u>0.001¢</u>	51%	<u>0.0074¢</u>
DoT	41%	<u>1.58¢</u>	63%	<u>1.20¢</u>	59%	<u>1.02¢</u>	58%	<u>0.84¢</u>	82%	<u>0.49¢</u>	50%	<u>0.13¢</u>
Router-R1	7%	<u>0.14¢</u>	2%	<u>0.15¢</u>	58%	<u>0.62¢</u>	47%	<u>9.78¢</u>	54%	<u>0.12¢</u>	38%	<u>0.12¢</u>
R2-Reasoner	38%	<u>1.16¢</u>	75%	<u>0.64¢</u>	76.5%	<u>0.08¢</u>	59.5%	<u>0.28¢</u>	83.5%	<u>0.042¢</u>	56.5%	<u>0.029¢</u>
Improvement	<u>↓9.52%</u>	<u>↓73.93%</u>	<u>↑10.29%</u>	<u>↓76.73%</u>	<u>↑21.43%</u>	<u>↓99.18%</u>	<u>↑2.59%</u>	<u>↓66.67%</u>	<u>↑1.83%</u>	<u>↓91.43%</u>	<u>↓4.24%</u>	<u>↓98.82%</u>

Table 1: Performance of R2-Reasoner and baselines on 6 benchmarks. C_{API} is averaged expense for each task, where API cost is measured in US dollar cents (¢). “-” appears in experiments where reasoning is conducted solely using local deployed SLMs without invoking the cloud-based LLMs. The highest reasoning accuracy is highlighted in bold. Results of the baseline with the highest Acc are underlined which will be used to compute the “Improvement” in the last row.

Selection and Deployment of LLMs For candidate LLMs to solve different subtasks, We select Qwen2.5-0.5B-instruct, Qwen2.5-1.5B-instruct, Qwen2.5-3B-instruct, Qwen2.5-7B-instruct, Qwen2.5-14B-instruct, Qwen2.5-32B-instruct, Qwen2.5-72B-instruct (Qwen et al., 2025), DeepSeek-V3 (DeepSeek-AI et al., 2025), gpt-4o (OpenAI, 2024) as the LLM pool. The ability of these LLMs increases following the order above. Among these models, Qwen2.5-0.5B-instruct, Qwen2.5-1.5B-instruct, Qwen2.5-3B-instruct, Qwen2.5-7B-instruct are fee free for being locally deployed, while the other cloud-based LLMs charges, and the price of the these LLMs also increases following the order above. For SFT and RL training on the task decomposer and subtask allocator, we select Qwen2.5-7B-instruct as the base model.

Evaluation For evaluation, we set two metrics: Acc and C_{API} , which represents our two main concerns in LLM reasoning. Acc measures the accuracy of our framework and the baselines on four benchmarks. C_{API} measures the average API cost for a single task, calculated in US dollar cents.

5.2 MAIN RESULTS

The comparison between our framework and the baselines in six benchmarks are shown in Table 1. We have highlighted in bold the highest accuracy results among the four baseline experiments on each benchmark, while the associated time costs and API costs are underlined. We compute the relative improvement of our results compared to the baseline with the highest accuracy. The experimental results demonstrate that our framework significantly reduces the API cost while retaining a comparable reasoning accuracy. The relative changes in accuracy compared to the highest baseline accuracy are: -9.52%, +10.29%, +21.43%, +2.59%, +1.83%, -4.24%. Even for P3, the decline in accuracy is still acceptable. The boost in accuracy on benchmark like MATH and SCAN validate the potential of our work in enhancing reasoning ability. Meanwhile, our framework achieves a tremendous reduction in API cost compared to the baseline with the highest accuracy, reaching averagely a decline of 84.46%. The accuracy of our framework on benchmarks like MATH and SCAN surpassing the CoT and ToT method shows the potential disadvantage of excessive reasoning. It usually happens in reasoning process conducted by LLMs of large scale, often deviates from the correct and suitable answer for a subtask because it automatically proceed with reflective or divergent thinking. We design several precise and exquisite prompts attempting to avoid the phenomenon.

5.3 ABLATION STUDY

To rigorously evaluate the contribution and of each stage, we report the performance metrics (Acc and C_{API}) of the Task Decomposer after each training stage, as summarized in Table 2. The table compares the base model, the model after supervised fine-tuning (SFT), and the final model after SFT combined with RL.

As observed, the SFT stage improves performance across all benchmarks compared to the base model. Importantly, the addition of the RL stage consistently further enhances both accuracy and cost efficiency on every task. For instance, accuracy increases by 5–8% on most benchmarks, while

Stages	P3		SCAN		MATH		CHAMP		CSQA		MuSiQue	
	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}
base	23.5%	0.314¢	14%	0.066¢	67%	0.150¢	50%	0.494¢	70.5%	0.147¢	43%	0.0226¢
w/ SFT	33%	2.027¢	68%	0.577¢	75.5%	0.079¢	58%	0.370¢	82%	0.056¢	51.5%	0.0301¢
w/ SFT+RL	38%	1.160¢	75%	0.636¢	76.5%	0.080¢	59.5%	0.280¢	83.5%	0.042¢	56.5%	0.0287¢

Table 2: Performance (Acc and C_{API}) after each training stage.

C_{API} is reduced or maintained at a comparable level. This consistent improvement demonstrates that the RL stage not only reliably enhances task performance but also stabilizes the routing decisions across tasks. Overall, these results strongly validate the effectiveness and robustness of our RL-based multi-stage training process.

5.4 GENERALIZATION TO NEWLY UNSEEN LLMs

To evaluate the generalization capability of the proposed R2-Reasoner, we conduct an additional experiment in which several models are replaced with alternatives of comparable capacity, without retraining the framework. Specifically, Qwen2.5-7B is replaced with GLM-4-9B-Chat (GLM et al., 2024), and DeepSeek-V3 with Kimi-K2-Instruct (Team et al., 2025). The results are summarized in Table 3.

Models	P3		SCAN		MATH		CHAMP		CSQA		MuSiQue	
	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}
Initial Pool	38%	1.160¢	75%	0.636¢	76.5%	0.080¢	59.5%	0.280¢	83.5%	0.042¢	56.5%	0.0287¢
Modified Pool	33.5%	1.278¢	75%	0.656¢	75%	0.105¢	51.5%	0.310¢	81.5%	0.060¢	51.5%	0.0438¢

Table 3: Experimental results of generalization capability of R2-Reasoner to new LLMs

As observed, the performance of our framework remains largely stable on SCAN, MATH and CSQA. Accuracy decreases by 11.8% on P3, 13% on CHAMP and 9% on MuSiQue, which can be attributed to differences in the reasoning capabilities of the replaced models. Meanwhile, C_{API} increases due to the higher API costs associated with the new models. Overall, these results indicate that the framework exhibits robust generalization to previously unseen LLMs. Importantly, the R2-Reasoner does not rely on any particular model; as long as the relative ordering of model capabilities is preserved, the router can maintain stable and reliable performance across different model pools.

5.5 TRADE-OFF BETWEEN REASONING COST AND ACCURACY

Our framework supports a flexible trade-off between accuracy and cost, enabling adaptation to different budget scenarios. By adjusting the routing threshold within our R2-Reasoner, we can dynamically balance performance and expenditure. As shown in Figure 3, when compared against the DoT and DataShunt baselines on the MATH and SCAN benchmarks, our method establishes a new Pareto frontier. The results clearly show that R2-Reasoner consistently achieves significantly higher accuracy for a given cost budget, or conversely, reaches a target accuracy at a substantially lower cost than both competing methods.

This remarkable efficiency is quantitatively demonstrated across both datasets. On the MATH benchmark, R2-Reasoner achieves over 70% accuracy for less than 0.08 cents, while the stronger baseline, DoT, requires approximately 6 cents to reach similar performance—a cost reduction of more than 75×. This advantage holds on the SCAN dataset, where our method reaches 60% accuracy for about 0.4 cents, a task that costs the DoT baseline approximately 5 cents. These results empirically prove that our routing mechanism enables highly effective and budget-aware reasoning, offering practical adaptability for diverse real-world deployment scenarios with varying budget constraints.

5.6 INFERENCE TIME COMPARISON ACROSS LLM ROUTERS

We conducted additional experiments under a consistent network environment to evaluate the end-to-end reasoning latency of our framework against several baseline methods. Each experiment was performed independently under identical conditions. All API calls were made sequentially in a single

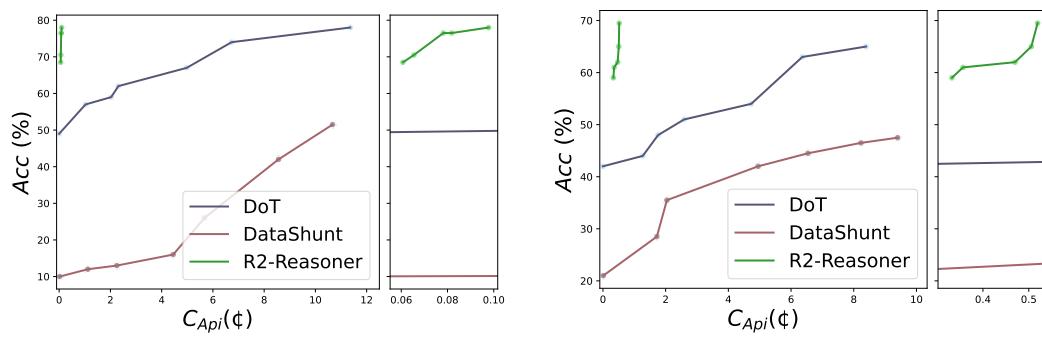


Figure 3: Acc-Cost trade-off curves on MATH (left) and SCAN (right). A magnified inset is provided to the right of the original sub-figure to more precisely illustrate the Pareto frontier of our method.

thread to eliminate concurrency-related interference and ensure that external factors did not distort the latency measurements. The reported results represent the average latency across all tasks in the benchmark, computed after completing full inference runs for every task. In each bar plot, the bar with the darkest color corresponds to our proposed method. The summarized results are presented in Figure 4.

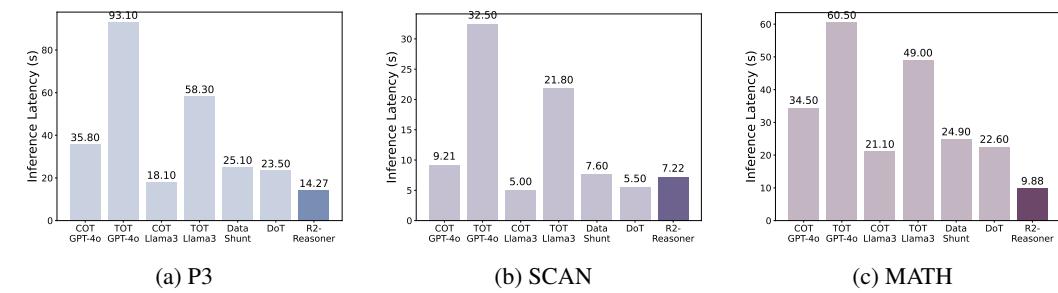


Figure 4: Inference latency comparison of different methods across three benchmarks.

The inference latency results demonstrate significant differences among the evaluated routing methods across the four benchmark tasks. Notably, R2-Reasoner consistently achieves the lowest or near-lowest latency in most cases. For instance, on P3, R2-Reasoner completes inference in 14.27 seconds, substantially faster than CoT and ToT configurations with both GPT-4o and LLaMA 3-8B models, which require between 18.1 and 93.1 seconds. Similar trends are observed on MATH and CSQA, where R2-Reasoner reduces inference time by more than 50% compared to the heaviest baselines (ToT).

On SCAN, R2-Reasoner incurs a slightly higher latency than CoT (LLaMA 3-8B), but it still remains considerably faster than the majority of other methods, including all GPT-4o-based baselines. This performance advantage can be attributed to the framework’s adaptive routing strategy, which prioritizes lightweight models for simpler instances and selectively invokes higher-capacity models only when necessary. As a result, R2-Reasoner achieves both time efficiency and cost efficiency, without compromising task performance. Overall, these results highlight the framework’s capability to perform fast and scalable reasoning across diverse benchmarks, demonstrating clear practical advantages over existing LLM routing methods.

6 CONCLUSION

In this work, we presented **R2-Reasoner**, a novel framework leveraging a reinforced Model Router to efficiently scale large language model reasoning by decomposing complex tasks and allocating subtasks to heterogeneous models. Our staged training pipeline, combining supervised fine-tuning with iterative reinforcement learning, enables adaptive, cost-effective collaboration among models. Looking forward, R2-Reasoner offers promising potential for real-world applications requiring scalable, resource-aware multi-model reasoning, such as complex decision-making systems and cloud computing platform.

486 REPRODUCIBILITY STATEMENT
487488 To ensure the reproducibility of our work, we provide all necessary resources and code used in this
489 paper. All benchmarks and models employed are fully open-source or publicly accessible, and no
490 privacy or copyright concerns are involved. All datasets and models are cited in Section 5.1.491 Our project code, including the implementation of the R2-Reasoner framework, training scripts,
492 and evaluation pipelines, is publicly available via the following anonymous link: https://anonymous.4open.science/r/R2_Reasoner.
493494 Additionally, the main paper, appendix C.2, C.1, and supplementary materials include detailed
495 descriptions of the experimental setup, hyperparameters, and evaluation protocols. Together with the
496 provided code, these materials allow other researchers to fully reproduce the results reported in this
497 work.
498499 REFERENCES
500501 Pranjal Aggarwal, Aman Madaan, Ankit Anand, Srividya Pranavi Potharaju, Swaroop Mishra, Pei
502 Zhou, Aditya Gupta, Dheeraj Rajagopal, Karthik Kappagantu, Yiming Yang, Shyam Upadhyay,
503 Manaal Faruqui, and Mausam. Automix: Automatically mixing language models, 2025. URL
504 <https://arxiv.org/abs/2310.12963>.505 Fenglong Cai, Dong Yuan, Zhe Yang, and Lizhen Cui. Edge-llm: A collaborative framework for
506 large language model serving in edge computing. In *2024 IEEE International Conference on Web
507 Services (ICWS)*, pp. 799–809. IEEE, 2024.508 Dong Chen, Yueling Zhuang, Shuo Zhang, Jinfeng Liu, Su Dong, and Siliang Tang. Data shunt:
509 Collaboration of small and large models for lower costs and better performance. In *Proceedings of
the AAAI Conference on Artificial Intelligence*, volume 38, pp. 11249–11257, 2024a.510 Guoxin Chen, Minpeng Liao, Chengxi Li, and Kai Fan. Step-level value preference optimization for
511 mathematical reasoning. *arXiv preprint arXiv:2406.10858*, 2024b.512 Shuhao Chen, Weisen Jiang, Baijiong Lin, James Kwok, and Yu Zhang. Routerdc: Query-based
513 router by dual contrastive learning for assembling large language models. *Advances in Neural
514 Information Processing Systems*, 37:66305–66328, 2024c.515 DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang
516 Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli
517 Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen,
518 Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding,
519 Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi
520 Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song,
521 Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang,
522 Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan
523 Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang,
524 Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi
525 Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li,
526 Shanghao Lu, Shangyan Zhou, Shanhua Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye,
527 Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang,
528 Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanjia Zhao, Wei An, Wen Liu,
529 Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyu Jin, Xianzu Wang,
530 Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha
531 Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu,
532 Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su,
533 Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong
534 Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng,
535 Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan
536 Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yuduan Wang, Yue
537 Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yan, Yuxiang Luo,
538

540 Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu,
 541 Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou,
 542 Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu
 543 Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan.
 544 Deepseek-v3 technical report, 2025. URL <https://arxiv.org/abs/2412.19437>.

545 Jasper Dekoninck, Maximilian Baader, and Martin Vechev. A unified approach to routing and
 546 cascading for llms. *arXiv preprint arXiv:2410.10347*, 2024.

548 Tao Feng, Yanzhen Shen, and Jiaxuan You. Graphrouter: A graph-based router for llm selections.
 549 *arXiv preprint arXiv:2410.03834*, 2024.

551 Team GLM, :, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego
 552 Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie
 553 Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Jingyu Sun, Juanzi Li, Lei Zhao, Lindong
 554 Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi
 555 Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao
 556 Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song,
 557 Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao
 558 Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang.
 559 Chatglm: A family of large language models from glm-130b to glm-4 all tools, 2024. URL
 560 <https://arxiv.org/abs/2406.12793>.

561 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu
 562 Zhang, Shirong Ma, Xiao Bi, et al. Deepseek-r1 incentivizes reasoning in llms through reinforce-
 563 ment learning. *Nature*, 645(8081):633–638, 2025.

564 Zixu Hao, Huiqiang Jiang, Shiqi Jiang, Ju Ren, and Ting Cao. Hybrid slm and llm for edge-cloud
 565 collaborative inference. In *Proceedings of the Workshop on Edge and Mobile Foundation Models*,
 566 pp. 36–41, 2024.

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

572 Brenden Lake and Marco Baroni. Generalization without systematicity: On the compositional skills
 573 of sequence-to-sequence recurrent networks. In *International conference on machine learning*, pp.
 574 2873–2882. PMLR, 2018.

576 En Li, Liekang Zeng, Zhi Zhou, and Xu Chen. Edge ai: On-demand accelerating deep neural network
 577 inference via edge computing. *IEEE Transactions on Wireless Communications*, 19(1):447–457,
 578 2019.

580 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 581 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step, 2023. URL
 582 <https://arxiv.org/abs/2305.20050>.

583 Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min
 584 Lin. Understanding r1-zero-like training: A critical perspective. *arXiv preprint arXiv:2503.20783*,
 585 2025.

587 Yujun Mao, Yoon Kim, and Yilun Zhou. Champ: A competition-level dataset for fine-grained
 588 analyses of llms’ mathematical reasoning capabilities. *arXiv preprint arXiv:2401.06961*, 2024.

590 OpenAI. Introducing openai o1, 2024. URL <https://openai.com/o1/>.

592 OpenAI. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>, 2024.

593 OpenAI. Introducing gpt-5. <https://openai.com/index/introducing-gpt-5/>, 2025.

594 Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan
 595 Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang,
 596 Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin
 597 Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi
 598 Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan,
 599 Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. Qwen2.5 technical report, 2025. URL
 600 <https://arxiv.org/abs/2412.15115>.

601 Tal Schuster, Ashwin Kalyan, Oleksandr Polozov, and Adam Tauman Kalai. Programming puzzles.
 602 *arXiv preprint arXiv:2106.05784*, 2021.

603

604 Chenyang Shao, Xinyuan Hu, Yutang Lin, and Fengli Xu. Division-of-thoughts: Harnessing hybrid
 605 language model synergy for efficient on-device agents. In *Proceedings of the ACM on Web
 606 Conference 2025*, pp. 1822–1833, 2025a.

607

608 Chenyang Shao, Xinyuan Hu, Yutang Lin, and Fengli Xu. Division-of-thoughts: Harnessing hybrid
 609 language model synergy for efficient on-device agents, 2025b. URL <https://arxiv.org/abs/2502.04392>.

610

611 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 612 Mingchuan Zhang, Y. K. Li, Y. Wu, and Daya Guo. Deepseekmath: Pushing the limits of
 613 mathematical reasoning in open language models, 2024. URL <https://arxiv.org/abs/2402.03300>.

614

615 Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling lilm test-time compute optimally
 616 can be more effective than scaling model parameters, 2024. URL <https://arxiv.org/abs/2408.03314>.

617

618 Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question
 619 answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*, 2018.

620

621 Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen,
 622 Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong,
 623 Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, Yichen Feng, Kelin Fu, Bofei Gao,
 624 Hongcheng Gao, Peizhong Gao, Tong Gao, Xinran Gu, Longyu Guan, Haiqing Guo, Jianhang
 625 Guo, Hao Hu, Xiaoru Hao, Tianhong He, Weiran He, Wenyang He, Chao Hong, Yangyang Hu,
 626 Zhenxing Hu, Weixiao Huang, Zhiqi Huang, Zihao Huang, Tao Jiang, Zhejun Jiang, Xinyi Jin,
 627 Yongsheng Kang, Guokun Lai, Cheng Li, Fang Li, Haoyang Li, Ming Li, Wentao Li, Yanhao
 628 Li, Yiwei Li, Zhaowei Li, Zheming Li, Hongzhan Lin, Xiaohan Lin, Zongyu Lin, Chengyin
 629 Liu, Chenyu Liu, Hongzhang Liu, Jingyuan Liu, Junqi Liu, Liang Liu, Shaowei Liu, T. Y. Liu,
 630 Tianwei Liu, Weizhou Liu, Yangyang Liu, Yibo Liu, Yiping Liu, Yue Liu, Zhengying Liu, Enzhe
 631 Lu, Lijun Lu, Shengling Ma, Xinyu Ma, Yingwei Ma, Shaoguang Mao, Jie Mei, Xin Men, Yibo
 632 Miao, Siyuan Pan, Yebo Peng, Ruoyu Qin, Bowen Qu, Zeyu Shang, Lidong Shi, Shengyuan
 633 Shi, Feifan Song, Jianlin Su, Zhengyuan Su, Xinjie Sun, Flood Sung, Heyi Tang, Jiawen Tao,
 634 Qifeng Teng, Chensi Wang, Dinglu Wang, Feng Wang, Haiming Wang, Jianzhou Wang, Jiaxing
 635 Wang, Jinhong Wang, Shengjie Wang, Shuyi Wang, Yao Wang, Yeqie Wang, Yiqin Wang, Yuxin
 636 Wang, Yuzhi Wang, Zhaoji Wang, Zhengtao Wang, Zhexu Wang, Chu Wei, Qianqian Wei, Wenhao
 637 Wu, Xingzhe Wu, Yuxin Wu, Chenjun Xiao, Xiaotong Xie, Weimin Xiong, Boyu Xu, Jing Xu,
 638 Jinjing Xu, L. H. Xu, Lin Xu, Suting Xu, Weixin Xu, Xinran Xu, Yangchuan Xu, Ziyao Xu, Junjie
 639 Yan, Yuzi Yan, Xiaofei Yang, Ying Yang, Zhen Yang, Zhilin Yang, Zonghan Yang, Haotian Yao,
 640 Xingcheng Yao, Wenjie Ye, Zhuorui Ye, Bohong Yin, Longhui Yu, Enming Yuan, Hongbang Yuan,
 641 Mengjie Yuan, Haobing Zhan, Dehao Zhang, Hao Zhang, Wanlu Zhang, Xiaobin Zhang, Yangkun
 642 Zhang, Yizhi Zhang, Yongting Zhang, Yu Zhang, Yutao Zhang, Yutong Zhang, Zheng Zhang,
 643 Haotian Zhao, Yikai Zhao, Huabin Zheng, Shaojie Zheng, Jianren Zhou, Xinyu Zhou, Zaida Zhou,
 644 Zhen Zhu, Weiyu Zhuang, and Xinxing Zu. Kimi k2: Open agentic intelligence, 2025. URL
 645 <https://arxiv.org/abs/2507.20534>.

646 Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. Musique: Multihop
 647 questions via single-hop question composition, 2022. URL <https://arxiv.org/abs/2108.00573>.

648 Peiyi Wang, Lei Li, Zhihong Shao, RX Xu, Damai Dai, Yifei Li, Deli Chen, Y Wu, and Zhifang
 649 Sui. Math-shepherd: A label-free step-by-step verifier for llms in mathematical reasoning. *arXiv*
 650 *preprint arXiv:2312.08935*, 2023.

651

652 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 653 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 654 *Neural Information Processing Systems*, 35:24824–24837, 2022.

655 Lilian Weng. Reward hacking in reinforcement learning. *lilianweng.github.io*, Nov 2024. URL
 656 <https://lilianweng.github.io/posts/2024-11-28-reward-hacking/>.

657

658 Noam Wies, Yoav Levine, and Amnon Shashua. Sub-task decomposition enables learning in sequence
 659 to sequence tasks, 2023. URL <https://arxiv.org/abs/2204.02892>.

660 Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws: An
 661 empirical analysis of compute-optimal inference for problem-solving with language models. *arXiv*
 662 *preprint arXiv:2408.00724*, 2024.

663

664 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik
 665 Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *arXiv*
 666 *preprint arXiv:2305.10601*, 2023.

667

668 Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong
 669 Liu, Lingjun Liu, Xin Liu, et al. Dapo: An open-source llm reinforcement learning system at scale.
 670 *arXiv preprint arXiv:2503.14476*, 2025.

671

672 Haozhen Zhang, Tao Feng, and Jiaxuan You. Router-r1: Teaching llms multi-round routing and
 673 aggregation via reinforcement learning. *arXiv preprint arXiv:2506.09033*, 2025.

674

675 Mingjin Zhang, Jiannong Cao, Xiaoming Shen, and Zeyang Cui. Edgeshard: Efficient llm inference
 676 via collaborative edge computing. *arXiv preprint arXiv:2405.14371*, 2024.

677

678 Huaixiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed H Chi, Quoc V Le, and
 679 Denny Zhou. Take a step back: Evoking reasoning via abstraction in large language models. *arXiv*
 680 *preprint arXiv:2310.06117*, 2023.

681

682 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans,
 683 Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning
 684 in large language models. *arXiv preprint arXiv:2205.10625*, 2022.

685

686 Wang Zhu, Jesse Thomason, and Robin Jia. Chain-of-questions training with latent answers for
 687 robust multistep question answering. *arXiv preprint arXiv:2305.14901*, 2023.

688

689 Yuxin Zuo, Kaiyan Zhang, Shang Qu, Li Sheng, Xuekai Zhu, Biqing Qi, Youbang Sun, Ganqu
 690 Cui, Ning Ding, and Bowen Zhou. Ttrl: Test-time reinforcement learning. *arXiv preprint*
 691 *arXiv:2504.16084*, 2025.

692

693

694

695

696

697

698

699

700

701

702 A SUPPLEMENTARY EXPERIMENT RESULTS

704 A.1 PERFORMANCE IMPROVEMENT OF TASK DECOMPOSER

706 To evaluate how the two stages of SFT and RL training have improved the task decomposer, we test
 707 100 tasks on four benchmark (one from each category) and report results using two global metrics:
 708 C_d and the comprehensive *Score* (defined in Section 4.1). The C_d is calculated as the accuracy of the
 709 final answer obtained by allocating all subtasks generated from the current checkpoint to Llama3-8B,
 710 while the *Score* is computed following Equation 3. The comparison between the base model and our
 711 trained checkpoints is shown in Table 4. On average, SFT and RL jointly yield a 27% increase in C_d
 712 and a 6% reduction in *Score*. Across all benchmarks, SFT provides consistent improvements, while
 713 RL exhibits mild instability but still contributes overall gains. We attribute this instability to potential
 714 insufficiencies in the reward function design.

Model	P3		SCAN		MATH		CSQA	
	C_d	Score	C_d	Score	C_d	Score	C_d	Score
base	0.06	2200.51	0.38	1600.19	0.28	1311.71	0.69	1171.53
w/ SFT	0.10	1848.40	0.46	1557.43	0.31	1265.60	0.75	1161.46
w/ SFT+RL	0.10	1788.41	0.45	1508.50	0.34	1234.63	0.72	1201.73

721 Table 4: Performance improvement achieved of the Task Decomposer after multi-stage training.

723 Beyond these global metrics, we further analyze decomposition quality on three finer-grained dimensions: *Conciseness*, *Practicality*, and *Coherence*. These dimensions are operationalized as follows:
 724 *Conciseness*: measured by the number of subtasks generated. *Practicality*: measured by the token
 725 cost required for reasoning. *Coherence*: measured by the proportion of logically incoherent subtask
 726 pairs (as described in Section 4.1).

Benchmark	Conciseness	Practicality	Coherence
SCAN	3.00→2.9263	2197.04→2208.89	0.1459→0.1367
MATH	7.54→4.36	848.03→939.45	0.0364→0.0116

733 Table 5: Evaluation of decomposition quality before and after the training pipeline.

735 Fewer subtasks are generally preferred, as they directly reduce API cost and latency. An excessive
 736 number of subtasks may cause redundancy and confusion. Token consumption is ideally lower, since
 737 concise answers are desirable, though moderately longer reasoning chains may yield more thorough
 738 inference. For coherence, a smaller value is better, indicating stronger logical consistency among
 739 subtasks. As shown in Table 5, our multi-stage training significantly improves decomposition quality
 740 across these dimensions, further validating the effectiveness of our approach.

742 A.2 PERFORMANCE IMPROVEMENT OF SUBTASK ALLOCATOR

744 To measure how the 2 stages of SFT and RL training have improved the ability of subtask allocator,
 745 we test 100 tasks on each benchmark and set 2 metrics for evaluation: *Acc* and *MAE*. The *Acc* metric
 746 measures how many allocation samples are correct according to the labels in our allocation dataset.
 747 The *MAE* metrics is based on the LLM pool listed below: Qwen2.5-0.5B-instruct, Qwen2.5-1.5B-
 748 instruct, Qwen2.5-3B-instruct, Qwen2.5-7B-instruct, Qwen2.5-14B-instruct, Qwen2.5-32B-instruct,
 749 Qwen2.5-72B-instruct, DeepSeek-V3, gpt-4o. Starting from Qwen2.5-0.5B-instruct as model 0, we
 750 sequentially assign model indices from 0 to 8, making the size of the number align with the scale
 751 of the LLMs. We calculate the *MAE* between the prediction LLM ID and the label LLM ID. The
 752 *MAE* metric indicates the distance on the LLM map, providing a supplementary sign showing that
 753 even if the prediction is wrong, how close it is to the labelled correct answer. The comparison of the
 754 base model and our training checkpoint are shown in Table 6. In overall the SFT and RL method
 755 have achieved on average 121.29% increase on accuracy and 24.08% decrease on *MAE*. On all
 benchmarks, the SFT method shows significant improvement in both metrics. RL method is also
 slightly unstable but still further achieve an overall improvement on the base of SFT method. The

756 insufficiency of RL method’s effect may be because some inevitable reward hacking during the RL
 757 process.
 758

Model	P3		SCAN		MATH		CSQA	
	Acc	MAE	Acc	MAE	Acc	MAE	Acc	MAE
base	0.0923	3.0763	0.1138	3.1041	0.1016	2.5355	0.1773	2.5638
w/ SFT	0.2197	2.7762	0.2067	1.9107	0.2362	1.8685	0.3274	1.9419
w/ SFT+RL	0.2187	2.7862	0.2606	1.9361	0.2410	1.8603	0.3227	1.9834

765 Table 6: Performance improvement achieved of the Subtask Allocator after multi-stage training.
 766
 767

768 A.3 RL TRAINING REDUCES DEPENDENCE ON SFT DATA

770 To further examine the effectiveness of the RL stage, we conducted an additional experiment on the
 771 MATH dataset by deliberately reducing the amount of supervised fine-tuning (SFT) data. Specifically,
 772 the SFT training set was reduced by 50%, while the number of RL training epochs was doubled.

773 Under this setting, the model’s accuracy initially dropped by 23% immediately after SFT due to
 774 the reduced amount of annotated data. However, after applying RL, not only was this performance
 775 degradation fully recovered, but the reasoning accuracy was further improved by an additional 1.5%
 776 compared to the original full-data SFT baseline.

777 This result highlights a key advantage of our RL process: beyond improving reasoning ability, it
 778 substantially reduces dependence on large quantities of annotated data. In practice, this suggests that
 779 RL can serve as a scalable alternative when labeled resources are limited, making our approach more
 780 data-efficient and broadly applicable.

782 A.4 EXPLORING MODEL ENSEMBLING

784 Our proposed framework supports flexible extensions and adaptations to different scenarios. As an
 785 illustrative case, we evaluate its capability on the task of *model ensembling*. Model ensembling is a
 786 widely used strategy that combines the outputs of multiple models in order to improve robustness and
 787 potentially enhance accuracy. To enable ensembling within our framework, we design a voting-based
 788 mechanism: multiple models are assigned to the same subtask in parallel, and the final answer is
 789 determined via majority voting. This mechanism serves as a drop-in replacement for the single-model
 790 assignment in our allocator.

791 As shown in Table 7, ensembling improves the accuracy on MATH, but it does not provide consistent
 792 advantages on other benchmarks. Upon further analysis, we suspect that introducing additional
 793 models may also introduce misleading signals, which can interfere with the reasoning process of
 794 the framework. This experiments demonstrate that our framework not only supports rerouting after
 795 failure, but also generalizes to multi-model allocation for a single subtask when ensemble behavior is
 796 desired.

797 A.5 CLARIFICATIONS ON POTENTIAL BIAS IN CONSTRUCTED TRAINING DATASET

800 In the construction of our task decomposition dataset, we applied uniform evaluation metrics across
 801 different models, such as token counts, to ensure comparability. To mitigate potential biases arising
 802 from inherent differences in how models generate responses (e.g., varying token length distributions),
 803 we conducted additional experiments. Specifically, we measured the average token consumption of

Model	P3		SCAN		MATH		CSQA	
	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}	Acc	C_{API}
R2-Reasoner	38%	1.160¢	75%	0.640¢	76.5%	0.080¢	83.5%	0.042¢
R2-Reasoner w/ Ensembling	38%	2.577¢	54.5%	0.934¢	83%	0.222¢	81.5%	0.105¢

809 Table 7: Performance comparison of R2-Reasoner with and without model ensembling.

each model on two benchmark datasets and computed the mean and standard deviation across models. The results are summarized in Table 8.

Model Name	CSQA Token Num	MATH Token Num
Qwen2.5-0.5B	97.22	192.15
Qwen2.5-1.5B	82.75	105.37
Qwen2.5-3B	69.72	61.17
Qwen2.5-7B	58.99	48.72
Llama3-8B	58.34	64.42
Qwen2.5-14B	49.03	78.65
Qwen2.5-32B	52.98	116.10
Qwen2.5-72B	48.51	95.01
DeepSeek-V3	63.91	60.26
GPT-4o	57.28	71.10
Mean	63.87	89.29
Std. Dev.	15.55	42.07

Table 8: Average token consumption on CSQA and MATH.

The results show no substantial or systematic variation in token consumption across models, indicating that differences in token usage are not a major source of variability. Instead, the primary source of cost variation lies in the per-token inference cost, which is positively correlated with the model’s parameter scale.

B FURTHER SUPPLEMENTS TO THE METHODS AND FORMULAS

B.1 DETAILED FORMULATION OF THE TASK DECOMPOSER

Here, we provide a detailed formulation of the dataset construction process for the Task Decomposer (4.1). The Task Decomposer, denoted as $\mathcal{M}_{\text{decomp}}$, is responsible for transforming a complex input task T into a sequence of clearly defined and logically connected subtasks: $T \xrightarrow{\mathcal{M}_{\text{decomp}}} \{t^1, t^2, \dots, t^k\}$, where k is the number of subtasks. To systematically evaluate and select high-quality decompositions, we define three complementary metrics. **Conciseness** measures the number of subtasks k , balancing between over-fragmentation and overly coarse decomposition. **Practicality** estimates the computational cost by summing the token usage of all subtasks under a baseline evaluation model $\mathcal{M}_{\text{eval}}$:

$$\text{Practicality}(d) = \sum_{i=1}^k \text{Tokens}(t^i, \mathcal{M}_{\text{eval}}). \quad (2)$$

Coherence evaluates the logical flow by counting adjacent subtask pairs that lack meaningful connection, denoted as $\text{Coe}_{\text{pair}}(d)$. Lower values indicate better continuity.

These metrics are combined into an overall score for a candidate decomposition $d = \{t^i\}_{i=1}^k$:

$$\text{Score}(d) = w_c \cdot k + w_p \cdot \sum_{i=1}^k \text{Tokens}(t^i, \mathcal{M}_{\text{eval}}) + w_d \cdot \text{Coe}_{\text{pair}}(d), \quad (3)$$

where $w_c, w_p, w_d > 0$ are weighting coefficients. Lower scores correspond to higher-quality decompositions.

Additionally, a binary correctness signal $C(d) \in \{0, 1\}$ is determined by attempting to solve the original task using decomposition d with the evaluation model $\mathcal{M}_{\text{eval}}$. For each task T , we generate a set of candidate decompositions $\mathcal{S}_T = \{d_1, d_2, \dots, d_m\}$ and select the decomposition d^* that minimizes the score while satisfying correctness if possible:

$$d^* = \begin{cases} \arg \min_{d \in \mathcal{S}_T, C(d)=1} \text{Score}(d) & \text{if any } C(d) = 1, \\ \arg \min_{d \in \mathcal{S}_T} \text{Score}(d) & \text{otherwise.} \end{cases} \quad (4)$$

864 The collection of all (T, d^*) pairs forms the decomposition dataset $\mathcal{D}_{\text{decomp}}$.
 865

866 Finally, the Task Decomposer is fine-tuned on $\mathcal{D}_{\text{decomp}}$ using a standard cross-entropy loss:
 867

$$\mathcal{L}_{\text{decomp}} = - \sum_{(T, d^*) \in \mathcal{D}_{\text{decomp}}} \sum_i \log P_{\theta_{\text{decomp}}}(d_i^* | T), \quad (5)$$

870 where d_i^* denotes the i -th subtask in the target decomposition. This training ensures that $\mathcal{M}_{\text{decomp}}$
 871 consistently generates concise, practical, and coherent subtask sequences suitable for efficient reasoning.
 872

874 **B.2 GROUPED SEARCH STRATEGY FOR ALLOCATOR TRAINING**

876 Here, we provide the full details of the grouped search algorithm used to construct the allocation
 877 dataset $\mathcal{D}_{\text{alloc}}$ (4.2).
 878

879 **Formal Problem.** Given subtasks $\{t^i\}$ from $\mathcal{M}_{\text{decomp}}$ and a model pool $\mathcal{M}_{\text{pool}}$, the objective is to
 880 find an allocation scheme M_A^* that minimizes resource consumption while ensuring correctness:
 881

$$M_A^* = \arg \min_{M_A} \mathbb{E}[C_{\text{Api}}(M_A) + C_{\text{Time}}(M_A)] \quad \text{s.t.} \quad \text{Acc}(M_A) = 1. \quad (6)$$

884 **Granularity Expansion.** Each subtask t^i is labeled with a difficulty level based on α -quantile token
 885 probabilities:
 886

$$G(t^i) = \begin{cases} G_E & p(t^i) \geq \tau_{\text{diff1}}, \\ G_M & \tau_{\text{diff2}} < p(t^i) < \tau_{\text{diff1}}, \\ G_H & p(t^i) \leq \tau_{\text{diff2}}. \end{cases} \quad (7)$$

890 Simultaneously, models are grouped by capability:
 891

$$\mathcal{M}_{\text{pool}} = \mathbb{G}_{\mathcal{M}}^{\text{SLM}} \cup \mathbb{G}_{\mathcal{M}}^{\text{MLM}} \cup \mathbb{G}_{\mathcal{M}}^{\text{LLM}}. \quad (8)$$

894 An initial allocation $M_{A,0}$ maps each subtask to the medium-capacity model within the corresponding
 895 group.
 896

897 **Within-Group Refinement.** For each iteration j , the allocation $M_{A,j}$ is updated as:
 898

$$M_{A,j+1}(t^i) = \begin{cases} \text{smaller}(\mathbb{G}_{\mathcal{M}}^X) & \text{if } \text{Acc}(M_{A,j}) = 1, \\ \text{larger}(\mathbb{G}_{\mathcal{M}}^X) & \text{if } \text{Acc}(M_{A,j}) = 0, \end{cases} \quad (9)$$

901 where $X = G(t^i)$.
 902

903 **Cross-Group Adjustment.** If correctness cannot be achieved with within-group adjustments, inter-
 904 group changes are made:
 905

$$M_{A,j+1}(t^i) \in \mathbb{G}_{\mathcal{M}}^Y, \quad Y \neq X, \quad (10)$$

907 subject to available model capacities.
 908

909 **Termination.** The algorithm halts after at most $N_{\text{iter_alloc}} \leq 20$ iterations or when $\text{Acc}(M_{A,j}) = 1$
 910 with minimal resource usage. The resulting allocations $\{(t^i, M_A^*)\}$ populate $\mathcal{D}_{\text{alloc}}$.
 911

912 **Training Objective.** The allocator $\mathcal{M}_{\text{alloc}}$ is trained on $\mathcal{D}_{\text{alloc}}$ via supervised fine-tuning. The loss
 913 function is defined as:
 914

$$\mathcal{L}_{\text{alloc}} = - \sum_{(\{t^i\}, M_A^*) \in \mathcal{D}_{\text{alloc}}} \sum_i \log P_{\theta_{\text{alloc}}}(M_A^*(t^i) | t^i). \quad (11)$$

916 The algorithmic workflow of grouped search strategy for allocator training is illustrated in Algorithm
 917 1.

B.3 GRPO OBJECTIVE FOR CO-TRAINING

For completeness, we provide the full GRPO objective function used in our co-training phase (4.3). The general form is:

$$\mathcal{J}_{\text{GRPO}}(\theta_{RL}) = \mathbb{E}_{q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q)} \left[\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{k=1}^{|o_i|} \left\{ \min \left[\frac{\pi_{\theta_{RL}}(o_{i,k}|q, o_{i,<k})}{\pi_{\theta_{\text{old}}}(o_{i,k}|q, o_{i,<k})} \hat{A}_{i,k}, \right. \right. \right. \right. \\ \left. \left. \left. \left. \text{clip} \left(\frac{\pi_{\theta_{RL}}(o_{i,k}|q, o_{i,<k})}{\pi_{\theta_{\text{old}}}(o_{i,k}|q, o_{i,<k})}, 1 - \epsilon, 1 + \epsilon \right) \hat{A}_{i,k} \right] - \beta \mathbb{D}_{\text{KL}}[\pi_{\theta_{RL}}(\cdot|q, o_{i,<k}) || \pi_{\text{ref}}(\cdot|q, o_{i,<k})] \right\} \right] \quad (12)$$

Here: - θ_{RL} are the parameters being optimized. - $\pi_{\theta_{RL}}$ is the current policy, while $\pi_{\theta_{old}}$ is the policy used to generate trajectories. - $o_{i,k}$ is the k -th action in the i -th trajectory given context q . - $\hat{A}_{i,k}$ is the estimated advantage for that action. - The first term is a clipped surrogate objective (as in PPO), and the second term penalizes deviation from a reference policy π_{ref} , controlled by β .

This formulation is applied identically when training either $\mathcal{M}_{\text{decomp}}$ or $\mathcal{M}_{\text{alloc}}$, depending on which module is currently being updated.

C EXPERIMENT DETAILS

C.1 EXPERIMENTAL ENVIRONMENT AND TRAINING HYPERPARAMETERS

The hardware environment used for our experiments and the specific training hyperparameters are summarized in Table 9.

In addition to the hardware specifications and basic training parameters, we also set hyperparameters during dataset construction. During constructing the dataset for the Task Decomposer, we computed a weighted average over the three dimensions of task decomposition, which involves three

hyperparameters: w_c , w_p , and w_d . These three hyperparameters serve as weights for: (1) the total number of subtasks, (2) the total number of tokens used during inference, and (3) the coherence score, respectively. Empirically, these three components exhibit significantly different value ranges across a wide range of tasks. Specifically, our analysis shows that their average values are approximately 5.87 (number of subtasks), 676.59 (token count), and 0.1541 (coherence score). To ensure the comparability of these components during weighted aggregation, our hyperparameter selection strategy is based on normalizing them to a similar scale. Accordingly, we set $w_c = 100$, $w_p = 1$, and $w_d = 1000$, which balances their contributions in the combined scoring function.

To assess the sensitivity, we also conducted experiments using 10 distinct parameter settings during the data construction process. We found that the dataset quality is generally stable when the weight ratios stay within a reasonable balance (i.e., fluctuating within ± 30). Among the parameters, the first one w_c plays the dominant role, critically affecting the quality of decomposition and the complexity of subsequent reasoning, while the other two serve auxiliary roles. We plan to include a more thorough sensitivity analysis in the revised manuscript.

Module	Element	Detail
System	OS	Ubuntu 20.04.6 LTS
	CUDA	12.4
	Python	3.12.9
	Pytorch	2.6.0
	trl	0.17.0
	accelerate	1.6.0
	peft	0.15.1
	flash_attn	2.7.4.post1
Workflow	Device	2*NVIDIA A100 80G
	API	Siliconflow & Microsoft Azure
SFT	Mode	Lora
	Batch size	4, 8
	Number of epochs	2, 3
	Max token length	2048
	Lora rank	32, 64
	Optimizer	AdamW
	Learning rate	0.00002, 0.00003
RL Training	Algorithm	GRPO
	Number of Generation	4
	Batch size	1
	Global step	1024
	Max token length	2048
	Optimizer	AdamW
	Learning rate	0.0001, 0.00015

Table 9: Detailed Experimental Settings

C.2 DETAILS OF THE BENCHMARKS

We verify the effectiveness of our framework upon six open-source benchmarks. These benchmarks target four distinct aspects of the model’s capability, including:

- **(1) Program Synthesis:** We select P3 (Schuster et al., 2021) (Python Programming Puzzle) for evaluation. P3 defines each puzzle by a python program f and evaluate the concerned ability of program synthesis by checking if the candidate input provided by machines could make f return `True`. By a form comprehensible for both humans and machines, it emphasizes on the ability involved during coding process such as syntax correctness and algorithmic reasoning.

1026 For task decomposer, we randomly choose 1500 puzzles from the original P3 benchmark,
 1027 and filter out 1085 puzzles with their decomposition results which are valid for SFT training.
 1028 For subtask allocator, we originally randomly select 2000 puzzles and eventually filter out
 1029 1687 puzzles with a total number of over 12000 subtask allocation samples as the candidate
 1030 dataset. We choose 4000 subtask allocation samples for SFT training on the subtask allocator.
 1031

- **(2) Language-Driven Navigation:** We select SCAN (Lake & Baroni, 2018) (Simplified version of the CommAI Navigation) for this evaluation. This benchmark consists of a set of navigation commands with the corresponding action sequences. By instructing machines to convert the navigation commands in natural language into a sequence of actions and comparing the generated sequence sample with the label, it focuses on assessing the ability of logical navigation, including traversal, backward reasoning and anomaly detection.

1032 For task decomposer, we randomly choose 2814 commands out of the original SCAN
 1033 benchmark, and filter out 1180 commands with their decomposition results for the SFT
 1034 stage training. For subtask allocator, we originally randomly select 2000 commands and
 1035 obtain a set of 7708 sub-command allocation samples. We also select 4000 sub-command
 1036 allocation samples for training the subtask allocator.
 1037

- **(3) Solving Math Problems:** We select MATH (Hendrycks et al., 2021) and CHAMP (Mao et al., 2024) for this evaluation. The MATH benchmark consists of 12,500 challenging competition mathematics problems, while the CHAMP benchmark contains 270 diverse high school competition-level math problems. These two mainly involve LLM's conducting computation, memorizing mathematical knowledge and utilizing problem-solving techniques. Solving math problems has been universally acknowledged as a crucial aspect to measure LLM's reasoning ability.

1038 For task decomposer, we randomly choose 2044 math problems from the original MATH
 1039 benchmark, and use 1430 problems with their decomposition results for fine-tuning the task
 1040 decomposer. For subtask allocator, we first select 2000 original math problems from the
 1041 benchmark. After building our own allocation dataset, we obtain over 7000 sub-problem
 1042 allocation samples, and choose 4000 sub-problem allocation samples for training and
 1043 boosting the ability of subtask allocator.
 1044

- **(4) Commonsense Reasoning:** We select CSQA (Talmor et al., 2018) (CommonsenseQA) and MuSiQue (Trivedi et al., 2022) for this evaluation. These 2 benchmarks require a broader commonsense knowledge base for LLM. Considering the knowledge base varies as the scale of LLM varies, it is a suitable benchmark to test if different LLMs in our framework could collaborate and compose an integrated knowledge base in commonsense scenario.

1045 For task decomposer, we randomly select 2273 commonsense queries from the original
 1046 benchmark, and utilize 1591 out of the queries to finish the training process. For subtask
 1047 allocator, we obtain an original dataset of 1800 commonsense queries, and obtain over 5500
 1048 sub-problem allocation samples, and select nearly 4000 sub-problem allocation samples for
 1049 SFT training.
 1050

D DISCUSSIONS

D.1 ENHANCING LLM REASONING VIA REINFORCEMENT LEARNING

1056 The reasoning process of LLMs can be formulated as a partially observable Markov decision process
 1057 (POMDP), where context serves as the state, token generation as the action, and the objective is to
 1058 learn a policy maximizing cumulative reward. Since DeepSeek-R1 (Guo et al., 2025), reinforcement
 1059 learning has become central to enhancing LLM reasoning (Yu et al., 2025; Zuo et al., 2025; Liu
 1060 et al., 2025). Group Relative Policy Optimization (GRPO) (Shao et al., 2024) has recently gained
 1061 widespread attention as a leading RL algorithm: it evaluates batches of outputs, computes relative
 1062 advantages, and uniformly assigns rewards across tokens. Unlike actor-critic methods relying on value
 1063 estimators, GRPO avoids estimation bias, instability, and reward hacking (Weng, 2024), achieving
 1064 more stable and faithful optimization.
 1065

1080 D.2 ALTERNATIVE REWARD DESIGNS
1081

1082 We considered several alternative reward formulations beyond the outcome-based design adopted in
1083 our experiments. One natural idea is to provide step-wise (intermediate) rewards. However, in tasks
1084 such as 24-point arithmetic or multi-step mathematical reasoning, it is often difficult to accurately
1085 assess the quality of intermediate steps without knowledge of the final outcome. This makes step-level
1086 reward annotation or computation unreliable in practice.

1087 Another potential direction is to leverage Monte Carlo Tree Search (MCTS) to approximate inter-
1088 mediate rewards. While this strategy can, in principle, provide more informative supervision, it
1089 introduces substantial computational overhead and significantly increases the complexity of the data
1090 construction pipeline, thereby limiting its scalability.

1091 Inspired by the outcome-driven reward design used in DeepSeek R1, we ultimately adopted a final-
1092 outcome-based reward scheme. This approach achieves a good balance between effectiveness and
1093 efficiency, while remaining scalable to large-scale training. Our experimental results demonstrate that
1094 this reward formulation is both practical and effective for the considered reasoning tasks.

1095 D.3 BROADER IMPACTS
1096

1097 Our R2-Reasoner framework has the potential to significantly broaden the accessibility and appli-
1098 cability of advanced AI reasoning capabilities. By substantially reducing computational costs and
1099 latency associated with complex multi-step reasoning, it can democratize the use of powerful Large
1100 Language Models. This could enable smaller organizations, individual researchers, or developers
1101 with limited resources to leverage state-of-the-art reasoning techniques that are currently prohibitively
1102 expensive. In practical terms, this could spur innovation across various sectors. For instance, in
1103 education, it could power more sophisticated and responsive AI tutors capable of breaking down
1104 complex problems for students in a cost-effective manner. In scientific research, R2-Reasoner could
1105 facilitate more intricate automated hypothesis generation and experimental design by making deep
1106 reasoning chains more feasible. For enterprise applications, it could lead to the development of more
1107 intelligent and nuanced customer service bots, data analysis tools, or decision support systems that
1108 can handle complex queries without incurring excessive operational costs.

1109 Furthermore, the principle of dynamically allocating resources based on sub-task complexity could
1110 inspire more sustainable AI practices. By preferentially using smaller, more energy-efficient models
1111 for simpler tasks, the overall energy consumption and carbon footprint associated with large-scale AI
1112 deployments could be reduced. The framework also encourages the development and utilization of
1113 a more diverse ecosystem of language models, fostering innovation in both large and small model
1114 architectures. Ultimately, by making sophisticated reasoning more efficient and economical, R2-
1115 Reasoner can help unlock new applications and accelerate the integration of AI into various aspects
1116 of daily life and industry, fostering more intelligent and adaptive systems.

1117

1118 E PROMPTS
1119

1120 Below show how we construct our prompts for the four aspects of model’s reasoning capability, each
1121 aspect taking one benchmark’s prompt as an example:

1122

1123 E.1 PROGRAM SYNTHESIS: P3
1124

1125 Below is the prompt for decomposition data collection on benchmark P3:

1126

1127

1128

1129

1130

1131

1132

1133

You will be provided with a Programming Puzzle. The ultimate task is
to find an input that will make the program return True.
To better accomplish this task, now you need to break the puzzle
into multiple steps, preferably between 3 and 8 steps.

These steps are organized in a chain-like manner, in which the steps
are supposed to be solved following a certain order.

```

1134
1135     Meanwhile when writing each broken-down question step on a separate
1136         line, the order of the questions should be the order of how to
1137         solve these broken-down question steps.
1138
1139     4 examples are as follows:
1140
1141     Program 1:
1142     def sat(li: List[int], k=6):
1143         def prod(nums):
1144             ans = 1
1145             for i in nums:
1146                 ans *= i
1147             return ans
1148             return min(li) > 1 and len(li) == k and all((1 + prod(li[:i] +
1149                 li[i + 1:])) % li[i] == 0 for i in range(k))
1150
1151     Result 1 of decomposed steps:
1152     step 1: Understand the conditions required by the function.
1153     step 2: Choose the length of the list based on k.
1154     step 3: Generate potential elements for the list.
1155     step 4: Calculate the product of all other elements for each element
1156             in the list when i = 0 and add 1 to the product.
1157     step 5: Calculate the product of all other elements for each element
1158             in the list when i = 1 and add 1 to the product.
1159     step 6: Calculate the product of all other elements for each element
1160             in the list when i = 2 and add 1 to the product.
1161     step 7: Calculate the product of all other elements for each element
1162             in the list when i = 3 and add 1 to the product.
1163     step 8: Calculate the product of all other elements for each element
1164             in the list when i = 4 and add 1 to the product.
1165     step 9: Calculate the product of all other elements for each element
1166             in the list when i = 5 and add 1 to the product.
1167     step 10: Verify the divisibility condition for each element.
1168     step 11: Adjust the elements and repeat until a valid list is found.
1169     step 12: Confirm that the list meets all conditions.
1170
1171
1172     Program 2:
1173     def sat(indices: List[int], s="aeEm%%uIV0imR&xUvQvZf#1z4"):
1174         i, j = indices
1175         return s[i] == s[j] and 0 <= i < j < i + 3
1176
1177     Result 2 of decomposed steps:
1178     step 1: Understand there are two conditions need to fulfill for the
1179             input indices that i and j in the indices should meet first s[i]
1180             == s[j] and 0 <= i < j < i + 3
1181     step 2: Iterate through the string sin a group of 3 characters, s[n]
1182             s[n+1] s[n+2]
1183     step 3: Compare the three characters to see if any of two characters
1184             are the same.
1185     step 4: If identical strings are found, Count the index of both % in
1186             the string s; If no identical characters, move to the
1187             consecutive three characters.
1188     step 5: Write the index of two identical characters and yield the
1189             final answer of list indices.
1190
1191
1192     Program 3:
1193     def sat(path: List[int], weights=[{{1: 20, 2: 1}}, {{2: 2, 3: 5}},
1194         {{1: 10}}], bound=11):
1195         return path[0] == 0 and path[-1] == 1 and sum(weights[a][b] for
1196             a, b in zip(path, path[1:])) <= bound
1197
1198     Result 3 of decomposed steps:
1199     step 1: Create a list that fulfill the first constraint to have 0 at
1200             index 0.

```

```

1188
1189     step 2: Create a list that fulfill the second constraint to have 1 at
1190     last.
1191     step 3: Given that the sum of weights[a][b] for a, b in zip(path,
1192         path[1:])) <= bound, we need to find values in the list weights
1193         that is less than 11.
1194     step 4: First checking if combining step 1 and step 2 to be path
1195         could be the correct input by calculating sum(weights[a][b] for
1196             a, b in zip(path, path[1:])) <= bound
1197     step 5: If the previous step is not correct, then think about what
1198         could be the integer filling between 0 and 1.
1199     step 6: Eliminate the incorrect candidates.
1200     step 7: Fill in the number to the list of integer.
1201     step 8: Verify the if the new list will make the function return
1202         True.

1203 Program 4:
1204 "name": "LastLetters:3",
1205 def sat(y: List[bool], x=['ryxadec', 'pyfixotibujadyxe',
1206                         'mopubywewexi witethig 7', '!',
1207                         'jethi sed c', 'lotextusavufubyb',
1208                         'wuxesafetateextysima pebutextiawafufok',
1209                         'tuchonip', 'S',
1210                         'xyvovikofutex pylekazuquekedajota E',
1211                         'wik xofoxujegerigubo ?',
1212                         'gipimakude 1', 'O', '^',
1213                         'lakiqiuuvuhenugu vajyquy P',
1214                         '6', 'fezore', 'vabithin
1215                         textusichytilejocoke',
1216                         'B', 'lasuthasebuvy que &',
1217                         'mymanuzuzudyc thazufys y', '',
1218                         'gecohywelawu', 'wath']):
1219     assert len(x) == len(y)
1220     for s, b in zip(x, y):
1221         if len(s.split(" ")[-1]) == 1:
1222             assert b == s[-1].isalpha()
1223         else:
1224             assert not b
1225     return True
1226 Result 4 of decomposed steps:
1227 step 1: Determine the length of the list x to ensure y has the same
1228     length.
1229 step 2: Loop through the list x to check the last word of each
1230     string.
1231 step 3: Check if the last segment of the string in x (seperated by
1232     space) have length 1.
1233 step 4: If Step 3 meet, check if that character is alphabetical
1234     characters.
1235 step 5: If step 4 is true, then the boolean value in list y with
1236     corresponding index should also be True. If not, False.
1237 step 6: If Step 3 do not meet, the boolean value in list y with
1238     corresponding index should be False.
1239 step 7: The final result should a list of boolean values.

1240 Now here is the puzzle for you to decompose: {question}
1241 Requirements:
1. The steps broken down should preferably be between 3 to 8 steps.
2. Each step needs to be executable, have a clear meaning, or
1242     produce a meaningful result.

1243 Answer Format:
1244 The process of solving the problem can be divided into the following
1245     steps:

```

```

1242
1243 1. question step 1
1244 2. question step 2
1245 3. question step 3
1246 ...
1247

```

1248 Below is the prompt for solving subtasks sequentially on benchmark P3:

```

1249
1250
1251 You will be provided with a Programming Puzzle. Your task is to find
1252 an input that will make the program return True.
1253 Here is the puzzle:{puzzle}
1254
1255 The data type of your final answer should be {ans_type}.
1256 I have broken this puzzle down into many easier subtasks.
1257 Following the order of the subtasks and solving every subtask in
1258 sequence lead to finding out the correct input.
1259 I will assign you sub-tasks one by one, and provide the results of
1260 the previous sub-tasks as a reference for your reasoning.
1261 Please follow the sequence of our subtasks to find the correct input
1262 .
1263
1264 Now, the first several subtasks are already solved, these subtasks
1265 listed below following the sequence:{previous_tasks}.
1266
1267 Their answers are listed below, also following the sequence:{previous_moves}.
1268
1269 Now you need to solve the subtask: {Step_dict[str(cnt)]}.
1270
1271 Focus exclusively on solving the subtask.
1272 Your answer should be concise and directly address the core
1273 reasoning process.
1274 Avoid any unnecessary comments, greetings, or expressions of
1275 enthusiasm. Only provide the essential reasoning process and
1276 answer.
1277
1278 Please provide the answer to the subtask.
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

```

1277 Below is the prompt for synthesizing to obtain the final answer on benchmark P3:

```

1278
1279
1280 We are provided with a programming puzzle. Our task is to find an
1281 input that will make the program return True.
1282 Here is the puzzle:{puzzle}
1283 The data type of our correct input should be {ans_type}.
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

```

```

1296
1297 Even if a subtask answer contains a reasoning mistake or calculation
1298 error, you must still use it as given.
1299 Do not infer the correct answer based on correct reasoning steps if
1300 the computed result is incorrect.
1301 Your final synthesis should reflect the exact values and conclusions
1302 stated in the subtask answers, even if they are incorrect.
1303
1304 The final answer is the input that will make the program return True
1305 .
1306 Please give the input and just give the answer without any
1307 additional explanation or clarification.
1308 for example, if the final answer is 3, you are supposed to output 3.
1309 To output "the answer is 3" is forbidden.
1310 for example, if the final answer is [1,2,3], you are supposed to
1311 output [1,2,3]. To output ```python [1,2,3]``` is forbidden.

```

1312 Below is the prompt for using task decomposer to decompose the original puzzle into a sequence of
1313 subtasks on benchmark P3:

```

1314
1315 You will be provided with a Programming Puzzle. The ultimate task is
1316 to find an input that will make the program return True.
1317 To better accomplish this task, now you need to break the puzzle
1318 into multiple steps, preferably between 3 and 8 steps.
1319
1320 These steps are supposed to be solved in a chain-like manner
1321 following a certain order.
1322 Meanwhile when writing each broken-down step on a separate line, the
1323 order of the steps should be the order of solving these broken-
1324 down steps.
1325
1 examples is as follows:
1326 Programming Puzzle:
1327 def sat(indices: List[int], s="aeEm%%uIV0imR&xUvQvZf#1z4"):
1328     i, j = indices
1329     return s[i] == s[j] and 0 <= i < j < i + 3
1330
1331 Answer:
1332 The process of solving the programming puzzle can be divided into
1333 the following steps:
1: Understand there are two conditions need to fulfill for the input
1334 indices that i and j in the indices should meet first s[i] == s
1335 [j] and 0 <= i < j < i + 3
2: Iterate through the string sin a group of 3 characters, s[n] s[n
1336 +1] s[n+2]
3: Compare the three characters to see if any of two characters are
1337 the same.
4: If identical strings are found, Count the index of both % in the
1338 string s; If no identical characters, move to the consecutive
1339 three characters.
5: Write the index of two identical characters and yield the final
1340 answer of list indices.
1341
1342 Now here is the puzzle for you to decompose: {original_question}
1343 Requirements:
1. The steps broken down should preferably be between 3 to 8 steps.
1344 2. Each step needs to be executable, have a clear meaning, or
1345 produce a meaningful result.
1346
1347 Answer Format:
1348 The process of solving the programming puzzle can be divided into
1349 the following steps:

```

```

1350
1351 1. step 1
1352 2. step 2
1353 3. step 3
1354 ...
1355

```

1356 Below is the prompt for using subtask allocator to allocate candidate model to a certain subtask on
1357 benchmark P3:

```

1358
1359
1360 Now we have a programming puzzle.
1361 We need to find the correct input that will make the program return
1362 true.
1363 We decide to break this puzzle into subtasks.
1364 Now we have to solve a subtask, and there are 9 models that can be
1365 chosen to solve this subtask.
1366 These 9 models are:
1367 qwen2.5-0.5b, qwen2.5-1.5b, qwen2.5-3b, qwen2.5-7b, qwen2.5-14b,
1368 qwen2.5-32b, qwen2.5-72b, deepseek-V3, gpt-4o.
1369 We list these models in ascending order according to their
1370 capability and the difficulty levels of the subtasks they are
1371 suitable for.
1372 For example,
1373 qwen2.5-0.5b has the lowest capability thus is suitable for the
1374 easiest subtask;
1375 gpt-4o has the highest capability thus is suitable for the hardest
1376 subtask.
1377 Task: choose the most appropriate model from the list above to solve
1378 the given subtask.
1379 Output only the chosen model's name.
1380
1381
1382 E.2 LANGUAGE-DRIVEN NAVIGATION: SCAN
1383
1384 Below is the prompt for decomposition data collection on benchmark SCAN:
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403

```

E.2 LANGUAGE-DRIVEN NAVIGATION: SCAN

Below is the prompt for decomposition data collection on benchmark SCAN:

```

1385
1386 I will give you a piece of natural language command. I need you to
1387 decompose it to smaller commands.
1388
1389 8 examples are as follows:
1390
1391 Command: "look right after look twice"
1392 Result of decomposition: "look right after look twice" can be solved
1393 by: "look right", "look twice".
1394
1395 Command: "jump opposite right thrice and walk"
1396 Result of decomposition: "jump opposite right thrice" can be solved
1397 by: "jump opposite right", "jump opposite right thrice". "walk"
1398 can be solved by: "walk". So, "jump opposite right thrice and
1399 walk" can finally be solved by: "jump opposite right", "jump
1400 opposite right thrice", "walk".
1401
1402 Command: "run left twice and run right"
1403 Result of decomposition: "run left twice" can be solved by: "run
1404 left", "run left twice". "run right" can be solved by "run right"
1405 ". So, "run left twice and run right" can finally be solved by:
1406 "run left", "run left twice", "run right".

```

```

1404
1405
1406 Command: "run opposite right"
1407 Result of decomposition: "run opposite right" can finally be solved
1408 by "run opposite right".
1409
1410 Command: "look opposite right thrice after walk"
1411 Result of decomposition: "look opposite right thrice" can be solved
1412 by: "look opposite right", "look opposite right thrice". "walk"
1413 can be solved by "walk". So, "look opposite right thrice after
1414 walk" can finally be solved by: "look opposite right", "look
1415 opposite right thrice", "walk".
1416
1417 Command: "jump around right"
1418 Result of decomposition: "jump around right" can be solved by: "jump
1419 right", "jump around right". So, "jump around right" can
1420 finally be solved by: "jump right", "jump around right".
1421
1422 Command: "look around right thrice and walk"
1423 Result of decomposition: "look around right thrice" can be solved by
1424 : "look right", "look around right", "look around right thrice".
1425 "walk" can be solved by "walk". So, "look around right thrice
1426 and walk" can finally be solved by: "look right", "look around
1427 right", "look around right thrice", "walk".
1428
1429 Command: "turn right after run right thrice"
1430 Result of decomposition: "turn right" can be solved by: "turn right
1431 ". "run right thrice" can be solved by: "run right", "run right
1432 thrice". So, "turn right after run right thrice" can finally be
1433 solved by: "turn right", "run right", "run right thrice".
1434
1435
1436 Now the command is {question}, please decompose it into smaller
1437 commands like the examples.
1438 Answer Format: xxx can be solved by: xxx. xxx can be solved by xxxx.
1439 ... So, xxx can finally be solved by: "subcommand_0", "
1440 subcommand_1",...
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457

```

Below is the prompt for solving sub-commands sequentially on benchmark SCAN:

```

1439 There is a natural language instruction representing a sequence of
1440 actions. I need you to translate this sentence from natural
1441 language into a standardized meta-action sequence."
1442 Here is the instruction:{question}
1443 I have broken this instruction down into some smaller instructions.
1444 I will assign you sub-instructions one by one, and provide the
1445 results of the previous sub-instructions as a reference for your
1446 reasoning.
1447 Please organize your reasoning according to the combination and
1448 progression of actions.
1449 For your reference, 13 examples for translation together with the
1450 corresponding explanations are as follows:
1451
1452 Q: "turn left"
1453 A: "turn left" outputs "TURN LEFT".
1454
1455 Q: "turn right"
1456 A: "turn right" outputs "TURN RIGHT".
1457 Q: "jump left"

```

1458
 1459 A: The output of `jump left` concatenates: the output of `turn`
 1460 `left`, the output of `jump`. `turn left` outputs
 1461 `TURN LEFT`. `jump` outputs `JUMP`. So concatenating
 1462 the output of `turn left` and the output of `jump`
 1463 leads to `TURN LEFT + JUMP`. So the output of `jump`
 1464 `left` is `TURN LEFT + JUMP`.
 1465
 1466 Q: "run right"
 1467 A: The output of "run right" concatenates: the output of "turn right"
 1468 ", the output of "run". "turn right" outputs "TURN RIGHT". "run"
 1469 outputs "RUN". So concatenating the output of "turn right" and
 1470 the output of "run" leads to "TURN RIGHT" + "RUN". So the output
 1471 of "run right" is "TURN RIGHT" + "RUN".
 1472
 1473 Q: "look twice"
 1474 A: The output of "look twice" concatenates: the output of "look",
 1475 the output of "look". "look" outputs "LOOK". So repeating the
 1476 output of "look" two times leads to "LOOK" * 2. So the output of
 1477 "look twice" is "LOOK" * 2.
 1478
 1479 Q: "run and look twice"
 1480 A: The output of "run and look twice" concatenates: the output of "
 1481 run", the output of "look twice". "run" outputs "RUN". "look
 1482 twice" outputs "LOOK" * 2. So concatenating the output of "run"
 1483 and the output of "look twice" leads to "RUN" + "LOOK" * 2. So
 1484 the output of "run and look twice" is "RUN" + "LOOK" * 2.
 1485
 1486 Q: "jump right thrice"
 1487 A: The output of "jump right thrice" concatenates: the output of "
 1488 jump right", the output of "jump right", the output of "jump
 1489 right". "jump right" outputs "TURN RIGHT" + "JUMP". So repeating
 1490 the output of "jump right" three times leads to ("TURN RIGHT" +
 1491 "JUMP") * 3. So the output of "jump right thrice" is ("TURN
 1492 RIGHT" + "JUMP") * 3.
 1493
 1494 Q: "walk after run"
 1495 A: The output of "walk after run" concatenates: the output of "run",
 1496 the output of "walk". "run" outputs "RUN". "walk" outputs "WALK"
 1497 ". So concatenating the output of "run" and the output of "walk"
 1498 leads to "RUN" + "WALK". So the output of "walk after run" is "
 1499 RUN" + "WALK".
 1500
 1501 Q: "turn opposite left"
 1502 A: The output of "turn opposite left" concatenates: the output of "
 1503 turn left", the output of "turn left". "turn left" outputs "TURN
 1504 LEFT". So repeating the output of "turn left" twice leads to "
 1505 TURN LEFT" * 2. So the output of "turn opposite left" is "TURN
 1506 LEFT" * 2.
 1507
 1508 Q: "turn around left"
 1509 A: The output of "turn around left" concatenates: the output of "
 1510 turn left", the output of "turn left", the output of "turn left
 1511 ", the output of "turn left". "turn left" outputs "TURN LEFT".
 1512 So repeating the output of "turn left" four times leads to "TURN
 1513 LEFT" * 4. So the output of "turn around left" is "TURN LEFT" *
 1514 4. Q: "turn opposite right" A: The output of "turn opposite
 1515 right" concatenates: the output of "turn right", the output of "
 1516 turn right". "turn right" outputs "TURN RIGHT". So repeating the
 1517 output of "turn right" twice leads to "TURN RIGHT" * 2. So the
 1518 output of "turn opposite right" is "TURN RIGHT" * 2.
 1519
 1520 Q: "turn around right"

1512
 1513 A: The output of "turn around right" concatenates: the output of "turn right", the output of "turn right", the output of "turn right", the output of "turn right". "turn right" outputs "TURN RIGHT". So repeating the output of "turn right" four times leads to "TURN RIGHT" * 4. So the output of "turn around right" is "TURN RIGHT" * 4.
 1514
 1515
 1516
 1517
 1518
 1519 Q: "walk opposite left"
 1520 A: The output of "walk opposite left" concatenates: the output of "turn opposite left", the output of "walk". "turn opposite left" outputs "TURN LEFT" * 2. "walk" outputs "WALK". So concatenating the output of "turn opposite left" and the output of "walk" leads to "TURN LEFT" * 2 + "WALK". So the output of "walk opposite left" is "TURN LEFT" * 2 + "WALK".
 1521
 1522
 1523
 1524
 1525
 1526 Q: "walk around left"
 1527 A: The output of "walk around left" concatenates: the output of "walk left", the output of "walk left", the output of "walk left". "walk left" outputs "TURN LEFT" + "WALK". So repeating the output of "walk around left" four times leads to ("TURN LEFT" + "WALK") * 4. So the output of "walk around left" is ("TURN LEFT" + "WALK") * 4.
 1528
 1529
 1530
 1531
 1532 Please pay attention to the use of parentheses.
 1533
 1534
 1535 Now, the first several sub-instructions are already solved, these
 1536 sub-instructions are listed below following the sequence:{
 1537 previous_steps}.
 1538 Their answers are listed below, also following the sequence:{
 1539 previous_answs}.
 1540 Now you need to solve the sub-instruction: {Step_dict[str(cnt)]}.
 1541
 1542 Focus exclusively on solving the sub-instruction.
 1543 Your answer should be concise and directly address the core
 1544 reasoning process.
 1545 Avoid any unnecessary comments, greetings, or expressions of
 1546 enthusiasm. Only provide the essential reasoning process and
 1547 answer.
 1548 Please provide the answer to the sub-instruction.
 1549

1550 Below is the prompt for synthesizing to obtain the final answer on benchmark SCAN:
 1551

1552
 1553 There is a natural language instruction representing a sequence of
 1554 actions. I need you to translate this sentence from natural
 1555 language into a standardized meta-action sequence."
 1556 Here is the instruction:{problem}
 1557
 1558 For your reference, 13 examples for translation together with the
 1559 corresponding explanations are as follows:
 1560
 1561 Q: "turn left"
 1562 A: "turn left" outputs "TURN LEFT".
 1563
 1564 Q: "turn right"
 1565 A: "turn right" outputs "TURN RIGHT".
 1566

1566
 1567 Q: "jump left"
 1568 A: The output of jump left concatenates: the output of turn
 left, the output of jump . turn left outputs
 TURN LEFT . jump outputs JUMP . So concatenating
 the output of turn left and the output of jump
 leads to TURN LEFT + JUMP . So the output of jump
 left is TURN LEFT + JUMP .

1569
 1570
 1571
 1572
 1573
 1574 Q: "run right"
 1575 A: The output of "run right" concatenates: the output of "turn right"
 ", the output of "run". "turn right" outputs "TURN RIGHT". "run"
 outputs "RUN". So concatenating the output of "turn right" and
 the output of "run" leads to "TURN RIGHT" + "RUN". So the output
 of "run right" is "TURN RIGHT" + "RUN".

1576
 1577
 1578
 1579 Q: "look twice"
 1580 A: The output of "look twice" concatenates: the output of "look",
 the output of "look". "look" outputs "LOOK". So repeating the
 output of "look" two times leads to "LOOK" * 2. So the output of
 "look twice" is "LOOK" * 2.

1581
 1582
 1583
 1584 Q: "run and look twice"
 1585 A: The output of "run and look twice" concatenates: the output of "run",
 the output of "look twice". "run" outputs "RUN". "look
 twice" outputs "LOOK" * 2. So concatenating the output of "run"
 and the output of "look twice" leads to "RUN" + "LOOK" * 2. So
 the output of "run and look twice" is "RUN" + "LOOK" * 2.

1586
 1587
 1588
 1589
 1590 Q: "jump right thrice"
 1591 A: The output of "jump right thrice" concatenates: the output of "jump right",
 the output of "jump right", the output of "jump right". "jump right" outputs "TURN RIGHT" + JUMP . So repeating
 the output of "jump right" three times leads to ("TURN RIGHT" +
 "JUMP") * 3. So the output of "jump right thrice" is ("TURN
 RIGHT" + "JUMP") * 3.

1592
 1593
 1594
 1595
 1596
 1597 Q: "walk after run"
 1598 A: The output of "walk after run" concatenates: the output of "run",
 the output of "walk". "run" outputs "RUN". "walk" outputs "WALK". So concatenating the output of "run" and the output of "walk"
 leads to "RUN" + "WALK". So the output of "walk after run" is "RUN" + "WALK".

1599
 1600
 1601
 1602
 1603 Q: "turn opposite left"
 1604 A: The output of "turn opposite left" concatenates: the output of "turn left",
 the output of "turn left". "turn left" outputs "TURN LEFT". So repeating the output of "turn left" twice leads to "TURN
 LEFT" * 2. So the output of "turn opposite left" is "TURN
 LEFT" * 2.

1605
 1606
 1607
 1608
 1609 Q: "turn around left"
 1610 A: The output of "turn around left" concatenates: the output of "turn left",
 the output of "turn left", the output of "turn left", the output of "turn left". "turn left" outputs "TURN LEFT". So repeating
 the output of "turn left" four times leads to "TURN LEFT" * 4. So the output of "turn around left" is "TURN LEFT" * 4.
 Q: "turn opposite right" A: The output of "turn opposite right" concatenates: the output of "turn right",
 the output of "turn right". "turn right" outputs "TURN RIGHT". So repeating the output of "turn right" twice leads to "TURN
 RIGHT" * 2. So the output of "turn opposite right" is "TURN RIGHT" * 2.

1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618 Q: "turn around right"
 1619

1620
 1621 A: The output of "turn around right" concatenates: the output of "turn right", the output of "turn right", the output of "turn right", the output of "turn right". "turn right" outputs "TURN RIGHT". So repeating the output of "turn right" four times leads to "TURN RIGHT" * 4. So the output of "turn around right" is "TURN RIGHT" * 4.
 1622
 1623
 1624
 1625
 1626
 1627 Q: "walk opposite left"
 1628 A: The output of "walk opposite left" concatenates: the output of "turn opposite left", the output of "walk". "turn opposite left" outputs "TURN LEFT" * 2. "walk" outputs "WALK". So concatenating the output of "turn opposite left" and the output of "walk" leads to "TURN LEFT" * 2 + "WALK". So the output of "walk opposite left" is "TURN LEFT" * 2 + "WALK".
 1629
 1630
 1631
 1632
 1633 Q: "walk around left"
 1634 A: The output of "walk around left" concatenates: the output of "walk left", the output of "walk left", the output of "walk left". "walk left" outputs "TURN LEFT" + "WALK". So repeating the output of "walk around left" four times leads to ("TURN LEFT" + "WALK") * 4. So the output of "walk around left" is ("TURN LEFT" + "WALK") * 4.
 1635
 1636
 1637
 1638
 1639
 1640 Please pay attention to the use of parentheses.
 1641
 1642
 1643 Following the order of the sub-instructions and solving every sub-instruction in sequence lead to the final answer.
 1644 All the sub-instructions are listed in the order: {previous_tasks}
 1645 The answers to all the sub-instructions are listed in the same order
 1646 : {previous_answs}
 1647
 1648 We can synthesize the final answer based on all the answers to the sub-instructions.
 1649 You must synthesize the final answer strictly based on the provided answers to the sub-instructions, without performing any error correction or independent recalculations.
 1650 Even if a sub-instruction answer contains a reasoning mistake or calculation error, you must still use it as given.
 1651 Do not infer the correct answer based on correct reasoning steps if the computed result is incorrect.
 1652 Your final synthesis should reflect the exact values and conclusions stated in the sub-instruction answers, even if they are incorrect.
 1653
 1654
 1655
 1656
 1657
 1658
 1659
 1660 Please give the final action sequence without any additional explanation or clarification.
 1661
 1662
 1663 Below is the prompt for translating a pseudo action sequence expression into a sequence of actions on benchmark SCAN:
 1664
 1665
 1666
 1667 Now I have a pseudo action sequence expression with parentheses and multiplication. I need you to help me convert this into a sequence of actions without an operator sign.
 1668
 1669
 1670 6 examples are as follows:
 1671
 1672 Q: "JUMP" * 3
 1673 Rewrite: "JUMP" * 3
 1674 A: 1 JUMP 2 JUMP 3 JUMP

```

1674
1675
1676 Q: "RUN" * 4 * 2
1677 Rewrite: "RUN" * 8
1678 A: 1 RUN 2 RUN 3 RUN 4 RUN 5 RUN 6 RUN 7 RUN 8 RUN
1679
1680 Q: "TURN RIGHT" + "WALK"
1681 Rewrite: "TURN RIGHT" + "WALK"
1682 A: TURN RIGHT WALK
1683
1684 Q: ("TURN LEFT" + "LOOK") * 2 + "TURN LEFT" + "LOOK"
1685 Rewrite: ("TURN LEFT" + "LOOK") * 2 + "TURN LEFT" + "LOOK"
1686 A: 1 (TURN LEFT LOOK) 2 (TURN LEFT LOOK) TURN LEFT LOOK
1687
1688 Q: ("TURN RIGHT" * 2 + "JUMP") * 4
1689 Rewrite: ("TURN RIGHT" * 2 + "JUMP") * 4
1690 A: 1 (1 TURN RIGHT 2 TURN RIGHT JUMP) 2 (1 TURN RIGHT 2 TURN RIGHT
1691 JUMP) 3 (1 TURN RIGHT 2 TURN RIGHT JUMP) 4 (1 TURN RIGHT 2 TURN
1692 RIGHT JUMP)
1693
1694 Q: "TURN LEFT" * 2 + ("TURN RIGHT" + "WALK") * 4 * 2
1695 Rewrite: "TURN LEFT" * 2 + ("TURN RIGHT" + "WALK") * 8
1696 A: 1 TURN LEFT 2 TURN LEFT 1 (TURN RIGHT WALK) 2 (TURN RIGHT WALK) 3
1697 (TURN RIGHT WALK) 4 (TURN RIGHT WALK) 5 (TURN RIGHT WALK) 6 (
1698 TURN RIGHT WALK) 7 (TURN RIGHT WALK) 8 (TURN RIGHT WALK)
1699
1700 The pseudo action sequence to be converted is as follows: {sentence}
1701 Please change it to the action sequences.
1702 Please JUST answer the result.

```

Below is the prompt for using task decomposer to decompose the original command into a sequence of sub-commands on benchmark SCAN:

```

1703
1704 I will give you a piece of natural language command. I need you to
1705 decompose it to smaller commands.
1706
1707 3 examples are as follows:
1708
1709 Command: "look right after look twice"
1710 Answer:
1711 The given command can finally be solved by: "look right", "look
1712 twice".
1713
1714 Command: "jump opposite right thrice and walk"
1715 Answer:
1716 The given command can finally be solved by: "jump opposite right", "
1717 jump opposite right thrice", "walk".
1718
1719 Command: "run left twice and run right"
1720 Answer:
1721 The given command can finally be solved by: "run left", "run left
1722 twice", "run right".
1723
1724 Now the command is {original_question}, please decompose it into
1725 smaller commands like the examples.
1726 Answer Format:
1727 The given command can finally be solved by: "subcommand_0", "
1728 subcommand_1", ...

```

Below is the prompt for using subtask allocator to allocate candidate model to a certain sub-command on benchmark SCAN:

```

1728
1729 Now we have an original command.
1730 To conduct this command, we decide to break this command into
1731 subcommands.
1732 Now we have to conduct a subcommand, and there are 9 models that can
1733 be chosen to conduct this subcommand.
1734 These 9 models are:
1735 qwen2.5-0.5b, qwen2.5-1.5b, qwen2.5-3b, qwen2.5-7b, qwen2.5-14b,
1736 qwen2.5-32b, qwen2.5-72b, deepseek-V3, gpt-4o.
1737 We list these models in ascending order according to their
1738 capability and the difficulty levels of the subcommands they are
1739 suitable for.
1740 For example,
1741 qwen2.5-0.5b has the lowest capability thus is suitable for the
1742 easiest subcommand
1743 gpt-4o has the highest capability thus is suitable for the hardest
1744 subcommand.
1745 Task: choose the most appropriate model from the list above to
1746 conduct the given subcommand.
1747 Output only the chosen model's name.
1748
1749
1750 the original command: {original_problem}
1751 the subcommand: {subtask}

```

E.3 SOLVING MATH PROBLEMS: MATH

Below is the prompt for decomposition data collection on benchmark MATH:

```

1752
1753 I will now give you a math problem. The type of problem is {type}.
1754 Please break this math problem down into several easy-to-solve
1755 steps.
1756
1757 These steps are organized in a chain-like manner, in which the steps
1758 are supposed to be solved following a certain order.
1759 Meanwhile when writing each broken-down step, the order of the steps
1760 should be the order of how to solve these broken-down question
1761 steps.
1762
1763 1 examples are as follows:
1764 Question: Four years ago, Kody was only half as old as Mohamed. If
1765 Mohamed is currently twice 30 years old, how old is Kody
1766 currently?
1767 Answer: To solve the question "How old is Kody currently?", we need
1768 to know: "How old is Mohamed currently?", "How old was Mohamed
1769 four years ago?", "How old was Kody four years ago?".
1770
1771 Now the command is {question}, please decompose it into easy-to-
1772 solve steps like the examples.
1773 Answer Format: (Please write each broken-down question step on a
1774 separate line, starting with a number.)
1775 To solve the question "xxx", we need to know:
1776 "1. question step 1",
1777 "2. question step 2",
1778 "3. question step 3".
1779

```

Below is the prompt for solving sub-problems sequentially on benchmark MATH:

```

1782
1783 You are provided with a math problem. Your task is to solve it and
1784 give it an answer.
1785 Here is the problem:\n{problem}
1786 The question belongs to the type pf {question_type}.
1787
1788 I have broken this problem down into many easier subproblems.
1789 Following the order of the subproblems and solving every subproblem
1790 in sequence lead to the final answer.
1791
1792 Now, the first several subproblems are already solved, these
1793 subproblems are listed below following their order:{  

1794 previous_tasks}.
1795 Their answers are listed below, also following their order:{  

1796 previous_answs}.
1797 Now you need to solve the subproblem: {Step_dict[str(cnt)]}.
1798
1799 Focus exclusively on solving the subproblem.
1800 Your answer should be concise and directly address the core
1801 reasoning process.
1802 Avoid any unnecessary comments, greetings, or expressions of
1803 enthusiasm. Only provide the essential reasoning process and
1804 answer.
1805 Please provide the answer to the subproblem.
1806
1807
1808 Below is the prompt for synthesizing to obtain the final answer on benchmark MATH:
1809
1810
1811 We are provided with a math problem. Our task is to solve it and
1812 give it an answer.
1813 Here is the problem:\{problem}
1814 The question belongs to the type pf {question_type}.
1815
1816 I have broken this problem down into many easier subproblems.
1817 Following the order of the subproblems and solving every subproblem
1818 in sequence lead to the final answer.
1819 All the subproblems are listed in the order: {previous_tasks}
1820 The answers to all the subproblems are listed in the same order: {  

1821 previous_answs}
1822
1823 We can synthesize the final answer based on all the answers to the
1824 subproblems.
1825 You must synthesize the final answer strictly based on the provided
1826 answers to the subproblems, without performing any error
1827 correction or independent recalculations.
1828 Even if a subproblem answer contains a reasoning mistake or
1829 calculation error, you must still use it as given.
1830 Do not infer the correct answer based on correct reasoning steps if
1831 the computed result is incorrect.
1832 Your final synthesis should reflect the exact values and conclusions
1833 stated in the subproblem answers, even if they are incorrect.
1834
1835 Please give the final answer without any additional explanation or
clarification.

```

1836 Below is the prompt for judging if the final answer is correct on benchmark MATH:
 1837
 1838
 1839 Here is a math problem with a standard answer and a student's answer
 1840 . Please help me determine if the student's answer is correct.
 1841 Problem: {problem}
 1842 question type: {question_type}
 1843
 1844 Standard answer: {solution}
 1845
 1846 Answer: {final_anw}

1847 If the student's answer is correct, just output True; otherwise,
 1848 just output False.
 1849 No explanation is required.

1850
 1851 Below is the prompt for using task decomposer to decompose the original problem into a sequence of
 1852 sub-problems on benchmark MATH:
 1853

1854
 1855 I will now give you a math problem. Please break this math problem
 1856 down into several easy-to-solve steps.
 1857
 1858 These sub-problems are supposed to be solved in a chain-like manner
 1859 following a certain order.
 1860 When writing each broken-down sub-problem, the order of the sub-
 1861 problems should be the order of solving these broken-down sub-
 1862 problems.
 1863
 1864 1 examples is as follows:
 1865 Question: Four years ago, Kody was only half as old as Mohamed. If
 1866 Mohamed is currently twice 30 years old, how old is Kody
 1867 currently?
 1868 Answer:
 1869 To solve the given question, we need to know:
 1870 1. How old is Mohamed cuurently?
 1871 2. How old was Mohamed four years ago?
 1872 3. How old was Kody four years ago?

1873
 1874 Now the command is {original_question}, please decompose it into
 1875 easy-to-solve steps like the examples.
 1876 Answer Format: (Please write each broken-down question step on a
 1877 separate line, starting with a number.)
 1878 To solve the given question, we need to know:
 1879 1. question step 1
 1880 2. question step 2
 1881 3. question step 3
 1882 ...

1883
 1884 Below is the prompt for using subtask allocator to allocate candidate model to a certain sub-problem
 1885 on benchmark MATH:
 1886

1887
 1888 Now we have an original problem.
 1889 To solve this problem, we decide to break this problem into
 1886 subproblems.
 1887 Now we have to solve a subproblem, and there are 9 models that can
 1888 be chosen to solve this subproblem.
 1889 These 9 models are:

```

1890
1891 qwen2.5-0.5b, qwen2.5-1.5b, qwen2.5-3b, qwen2.5-7b, qwen2.5-14b,
1892 qwen2.5-32b, qwen2.5-72b, deepseek-V3, gpt-4o.\n\n
1893 We list these models in ascending order according to their
1894 capability and the difficulty levels of the subproblems they are
1895 suitable for.
1896 For example,
1897 qwen2.5-0.5b has the lowest capability thus is suitable for the
1898 easiest subproblem
1899 gpt-4o has the highest capability thus is suitable for the hardest
1900 subproblem.
1901 Task: choose the most appropriate model from the list above to solve
1902 the given subproblem.
1903 Output only the chosen model's name.
1904
1905
1906
1907 E.4 COMMONSENSE REASONING: CSQA
1908 Below is the prompt for decomposition data collection on benchmark CSQA:
1909
1910
1911 I have a single-choice question involving common sense reasoning
1912 that I want to solve. I hope you can break down the problem-
1913 solving process into several sub-problems. You can consider
1914 analyzing the question itself as well as the options.
1915 The number of sub-problems doesn't need to be too many; each sub-
1916 problem should have a clear meaning and purpose.
1917
1918 These sub-problems are organized in a chain-like manner, in which
1919 the sub-problems are supposed to be solved following a certain
1920 order.
1921 Meanwhile when writing each broken-down sub-problem, the order of
1922 the sub-problems should be the order of how to solve these
1923 broken-down sub-problems.
1924
1925 8 examples are as follows:
1926 Question 1:
1927 You can read a magazine while waiting for your transportation
1928 on rails to arrive?
1929 Choices 1:
1930 A. Train station, B. Bookstore, C. Newsstand, D. Waiting room, E.
1931 Airport
1932 Answer 1:
1933 1. What does "waiting for your transportation on rails" indicate
1934 about your current location?
1935 2. Which place in the options can accommodate you reading a magazine
1936 ?
1937 3. Which places that satisfy question 2 are near your current
1938 location?
1939
1940 Question 2:
1941 If I wanted to see a lizard in its natural habitat but I do not
1942 speak Spanish, where would I go?
1943 Choices 2:
1944 A. Utah, B. South America, C. New Hampshire, D. Japan, E. New
1945 Mexico
1946 Answer 2:
1947 1. Which places are natural habitats for lizards?
1948 2. Which places have Spanish as the primary language?

```

1944 3. Combine the answers from sub-question 1 and sub-question 2, among
 1945 the natural habitats for lizards, which places do not speak
 1946 Spanish?
 1947

1948 Question 3:
 1949 John was stuck in his house. He couldn't get out the door. He was
 1950 very frightened when the smoke detectors went off, but luckily
 1951 it was a false alarm. Why might he be stuck?
 1952 Choices 3:
 1953 A. fire, B. belong to, C. winter storm, D. face south, E. burn down
 1954 Answer 3:
 1. What are possible reasons for being stuck in a house?
 1955 2. Which options are related to situations that might cause a person
 1956 to be stuck?
 1957 3. Why might these specific conditions make it difficult to leave
 1958 the house?
 1959

Question 4:
 1960 John was stuck in his house. He couldn't get out the door. He was
 1961 very frightened when the smoke detectors went off, but luckily
 1962 it was a false alarm. Why might he be stuck?
 1963 Choices 4:
 1964 A. fire, B. belong to, C. winter storm, D. face south, E. burn down
 1965 Answer 4:
 1. What are possible reasons for being stuck in a house?
 1966 2. Which options are related to situations that might cause a person
 1967 to be stuck?
 1968 3. Why might these specific conditions make it difficult to leave
 1969 the house?
 1970

Question 5:
 1971 When looking for a non-perishable food in your house, you'll often
 1972 go look in the?
 1973 Choices 5:
 1974 A. Stove, B. Table, C. Plate, D. Jar, E. Pantry
 1975 Answer 5:
 1. What is non-perishable food?
 1976 2. Where are non-perishable foods commonly stored in a household?
 1977 3. Which of the options (stove, table, plate, jar, pantry) is the
 1978 most logical place for storing non-perishable food?
 1979

Question 6:
 1980 What must elementary school students do when they are going into
 1981 class?
 1982 Choices 6:
 1983 A. Think for himself, B. Answer question, C. Wait in line, D. Speak
 1984 a foreign language, E. Cross road
 1985 Answer 6:
 1. What do elementary school students typically do before entering a
 1986 classroom?
 2. Which actions among the options are related to classroom entry
 1987 procedures?
 3. Why might students perform this action before entering the
 1988 classroom?
 1989

Question 7:
 1990 After eating dinner, having plenty to eat, and exercising, what is
 1991 likely to happen?
 1992 Choices 7:
 1993 A. Become tired, B. Indigestion, C. Flatulence, D. Become
 1994 intoxicated, E. Become full
 1995 Answer 7:
 1. What happens to the body after eating a large meal?
 1996

1998
 1999 2. What are common effects of exercising after eating?
 2000 3. Which of the options (become tired, indigestion, flatulence,
 2001 become intoxicated, become full) best matches the expected
 2002 outcome of eating a large meal followed by exercise?
 2003
 2004 Question 8:
 2005 He didn't like the walk up, but living on the top floor meant there
 2006 was nobody above him in the what?
 2007 Choices 8:
 2008 A. Apartment building, B. Tall building, C. Go down, D. Garden, E.
 2009 Office building
 2010 Answer 8:
 2011 1. What does "walk up" suggest about the type of building?
 2012 2. What kind of building would have a "top floor" and residents
 2013 living above or below each other?
 2014 3. Which option (apartment building, tall building, go down, garden,
 2015 office building) best describes a place where living on the top
 2016 floor would mean no one lives above?
 2017
 2018 Now the question is {question}, the options are: {options}, please
 2019 decompose it into sub-problems.
 2020 Answer Format: (Please write each broken-down sub-problem on a
 2021 separate line, starting with a number.)
 2022 To solve the question "xxx", we need to clarify / solve:
 2023 "1. sub-problem 1",
 2024 "2. sub-problem 2",
 2025 "3. sub-problem 3".
 2026
 2027

2028 Below is the prompt for solving sub-problems sequentially on benchmark CSQA:

2029
 2030
 2031
 2032 There is a single-choice question involving common sense reasoning.
 2033 I need you to solve it and give the right answer.
 2034 Here is the question:{problem}
 2035 Here are the options:{options}
 2036
 2037 I have broken this common sense reasoning question down into several
 2038 smaller subproblems.
 2039 Following the order of the subproblems and solving every subproblem
 2040 in sequence lead to the final answer.
 2041
 2042
 2043 Now, the first several subproblems are already solved, these
 2044 subproblems are listed below following their order:{
 2045 previous_tasks}.
 2046 Their answers are listed below, also following their order:{
 2047 previous_answs}.
 2048 Now you need to solve the subproblem: {Step_dict[str(cnt)]}.
 2049
 2050 Focus exclusively on solving the subproblem.
 2051 Your answer should be concise and directly address the core
 2052 reasoning process.
 2053 Avoid any unnecessary comments, greetings, or expressions of
 2054 enthusiasm. Only provide the essential reasoning process and
 2055 answer.
 2056
 2057 Please provide the answer to the subproblem.

2058 Below is the prompt for synthesizing to obtain the final answer on benchmark CSQA:

2052
 2053 There is a single-choice question involving common sense reasoning.
 2054 I need you to solve it and give the right answer.
 2055 Here is the question:{problem}
 2056 Here are the options:{options}
 2057
 2058 I have broken this common sense reasoning question down into several
 2059 smaller subproblems.
 2060 Following the order of the subproblems and solving every subproblem
 2061 in sequence lead to the final answer.
 2062 All the subproblems are listed in the order: {previous_tasks}
 2063 The answers to all the subproblems are listed in the same order: {
 2064 previous_answs}
 2065
 2066 We can synthesize the final answer based on all the answers to the
 2067 subproblems, and finally choose the letter of the correct option
 2068 .
 2069 You must synthesize the final answer strictly based on the provided
 2070 answers to the subproblems, without performing any error
 2071 correction or independent recalculations.
 2072 Even if a subproblem answer contains a reasoning mistake or
 2073 calculation error, you must still use it as given.
 2074 Do not infer the correct answer based on correct reasoning steps if
 2075 the computed result is incorrect.
 2076 Your final synthesis should reflect the exact values and conclusions
 2077 stated in the subproblem answers, even if they are incorrect.
 2078
 2079 Please provide only the letter of the option, without any additional
 2080 explanation or description.

2081 Below is the prompt for using task decomposer to decompose the original problem into a sequence of
 2082 sub-problems on benchmark CSQA:
 2083

2084
 2085 I have a single-choice question involving common sense reasoning
 2086 that I want to solve. I hope you can break down the problem-
 2087 solving process into several sub-problems. You can consider
 2088 analyzing the question itself as well as the options.
 2089 The number of sub-problems doesn't need to be too many; each sub-
 2090 problem should have a clear meaning and purpose.
 2091
 2092 These sub-problems are supposed to be solved in a chain-like manner
 2093 following a certain order.
 2094 When writing each broken-down sub-problem, the order of the sub-
 2095 problems should be the order of solving these broken-down sub-
 2096 problems.
 2097
 2098 1 example is as follows:
 2099 Question 1:
 2100 You can read a magazine while waiting for your transportation
 2101 on rails to arrive?
 2102 Choices 1:
 2103 A. Train station, B. Bookstore, C. Newsstand, D. Waiting room, E.
 2104 Airport
 2105 Answer 1:
 2106 To solve the given question, we need to clarify / solve:
 2107 1. What does "waiting for your transportation on rails" indicate
 2108 about your current location?

```

2106
2107 2. Which place in the options can accommodate you reading a magazine
2108   ?
2109 3. Which places that satisfy question 2 are near your current
2110   location?
2111
2112 Now the question is {original_question}, the options are: {options},
2113   please decompose it into sub-problems.
2114 Answer Format: (Please write each broken-down sub-problem on a
2115   separate line, starting with a number.)
2116 To solve the given question, we need to clarify / solve:
2117 1. sub-problem 1,
2118 2. sub-problem 2,
2119 3. sub-problem 3.
2120 ...
2121
2122 Below is the prompt for using subtask allocator to allocate candidate model to a certain sub-problem
2123 on benchmark CSQA:
2124
2125 Now we have an original problem.
2126 To solve this problem, we decide to break this problem into
2127   subproblems.
2128 Now we have to solve a subproblem, and there are 9 models that can
2129   be chosen to solve this subproblem.
2130 These 9 models are:
2131 qwen2.5-0.5b, qwen2.5-1.5b, qwen2.5-3b, qwen2.5-7b, qwen2.5-14b,
2132   qwen2.5-32b, qwen2.5-72b, deepseek-V3, gpt-4o.\n\n
2133 We list these models in ascending order according to their
2134   capability and the difficulty levels of the subproblems they are
2135   suitable for.
2136 For example,
2137 qwen2.5-0.5b has the lowest capability thus is suitable for the
2138   easiest subproblem
2139 gpt-4o has the highest capability thus is suitable for the hardest
2140   subproblem.
2141 Task: choose the most appropriate model from the list above to solve
2142   the given subproblem.
2143 Output only the chosen model's name.
2144
2145 the original problem: {original_problem}
2146 the subproblem: {subtask}
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159

```

You may include other additional sections here.