

Graph Reasoning Paradigm: Structured and Symbolic Reasoning with Topology-Aware Reinforcement Learning for Large Language Models

Anonymous ACL submission

Abstract

Long Chain-of-Thought (LCoT), achieved by Reinforcement Learning with Verifiable Rewards (RLVR), has proven effective in enhancing the reasoning capabilities of Large Language Models (LLMs). However, reasoning in current LLMs is primarily generated as plain text, where performing semantic evaluation on such unstructured data creates a computational bottleneck during training. Despite RLVR-based optimization, existing methods still suffer from coarse-grained supervision, reward hacking, high training costs, and poor generalization. To address these issues, we propose the Graph Reasoning Paradigm (GRP), which realizes structured and symbolic reasoning, implemented via graph-structured representations with step-level cognitive labels. Building upon GRP, we further design Process-Aware Stratified Clipping Group Relative Policy Optimization (PASC-GRPO), which leverages structured evaluation to replace semantic evaluation, achieves process-aware verification through graph-structured outcome rewards, and mitigates reward hacking via stratified clipping advantage estimation. Experiments demonstrate significant improvements across mathematical reasoning and code generation tasks. Data, models, and code will be released later.

1 Introduction

Long chain-of-thought (LCoT) has been proven effective for eliciting the reasoning capabilities of Large Language Models (LLMs), particularly when combined with reinforcement learning with verifiable rewards (RLVR) (El-Kishky, 2024; Chen et al., 2025; He et al., 2025). Such approaches have been applied to mathematical reasoning (Zhou et al., 2024; Moshkov et al., 2025) and code generation (El-Kishky et al., 2025; Yu et al., 2025), which are representative complex reasoning tasks and have become common benchmarks for evaluating the reasoning abilities of LLMs.

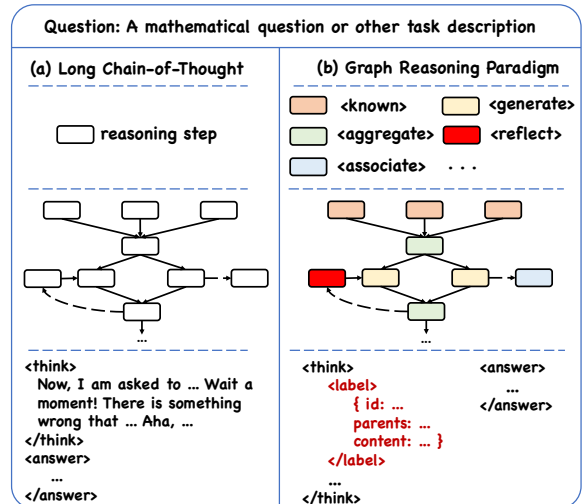


Figure 1: Comparison between traditional Long Chain-of-Thought and Graph Reasoning Paradigm.

However, current LLMs predominantly produce plain-text outputs. Although complex reasoning behaviors such as reflection and Aha moments have emerged (Yang et al., 2025b), the reliability of the reasoning process remains limited (Shojaee et al., 2025). Recently, several reinforcement learning (RL) approaches have attempted to address this issue. However, outcome-based reward methods provide only coarse-grained supervision, offering little semantic control over the reasoning process (DeepSeek-AI et al., 2025). Although process-based reward methods offer semantic control over the reasoning process, they often suffer from high training costs and poor generalization (Yuan et al., 2024). In summary, semantic evaluation is inherently difficult. Since LLMs output plain text, we are forced to perform semantic assessments, which creates a vicious cycle.

Inspired by the structured cognitive mechanisms of human cognition (George et al., 2021), we propose to replace plain-text with structured and symbolic representations, and to replace semantic eval-

uation with structure-based evaluation, thereby enabling low-cost and fine-grained control over the reasoning process when computing outcome rewards. Notably, increasing reliance on process-based rewards may encourage models to over-optimize intermediate reasoning quality at the expense of correct final answers, giving rise to reward hacking (Denison et al., 2024). This issue fundamentally reflects the limitations of conventional advantage estimation methods.

In this paper, we propose a **Graph Reasoning Paradigm (GRP)**, which realizes the structured and symbolic reasoning process using graph-structured representations. As shown in Figure 1, the reasoning process in this paradigm is explicitly organized into a structured form, where each reasoning step is annotated with a specific cognitive label. These labels correspond to different cognitive operations, including *known*, *generate*, *aggregate*, *reflect*, *refine*, *reverse*, and *associate*. This structured representation transforms reasoning from unstructured text into an explicit graph, improving interpretability, and enabling systematic evaluation of reasoning processes. Based on this paradigm, we construct over 40.3k graph-structured mathematical reasoning chains and 12k graph-structured code generation chains with step-level labels, and perform supervised fine-tuning (SFT) to train LLMs to internalize graph-structured reasoning.

Building on the structured and symbolic outputs, we propose **PASC-GRPO**, a **Process-Aware Stratified Clipping** extension of **Group Relative Policy Optimization** to further exert finer control over the reasoning process. We design a set of graph-based outcome rewards that evaluate the quality of the reasoning process using structural properties such as label validity, reachability, connectivity, informative subgraph, and reverse search consistency. These rewards do not rely on value models for semantic evaluation, leading to better generalization and improved training efficiency. To mitigate reward hacking introduced by multiple rewards, we further propose stratified clipping advantage estimation, which separately normalizes rewards within correct and incorrect groups, ensuring stable and reliable optimization.

Experiments across mathematical reasoning and code generation benchmarks demonstrate that our method consistently outperforms strong baselines, particularly on competition-level tasks. Furthermore, it effectively mitigates reward hacking and reduces reasoning length, improving both accuracy

and inference efficiency. In summary, our contributions are threefold:

- We propose a **Graph Reasoning Paradigm** that enables structured and symbolic reasoning processes. We construct over **52k** graph-structured reasoning chains with step-level cognitive labels for SFT.
- We propose **PASC-GRPO**, a reinforcement learning method that leverages graph-structured outcome rewards and stratified clipping advantage estimation to improve reasoning quality while mitigating reward hacking.
- We empirically demonstrate that our approach significantly improves the reasoning performance of LLMs, providing evidence for the effectiveness of graph reasoning paradigm.

2 Related Work

Structured cognition. Structured reasoning frameworks organize the reasoning process into explicit and manageable representations. Examples include Tree-of-Thoughts (Yao et al., 2023), Graph-of-Thoughts (Besta et al., 2023), and cognitive-architecture-inspired models (Sumers et al., 2023). However, these approaches primarily focus on prompting strategies, and rely on semantic evaluation or manual annotation.

RL-enhanced Reasoning. Reinforcement learning has been widely applied to improve the reasoning capabilities of LLMs (Xu et al., 2025; Zhang et al., 2025; Yuan et al., 2024). However, these methods often rely on coarse reward signals, making fine-grained process control and evaluation difficult, and increasing the risk of reward hacking (Denison et al., 2024). Detailed technical background is provided in Appendix A.

3 Methodology

As illustrated in Figure 2, our method consists of two main components:

- **Graph Reasoning Paradigm**, which enables LLMs to produce symbolized, interpretable, verifiable, and accurate graph-structured reasoning through supervised fine-tuning;
- **PASC-GRPO**, a reinforcement learning framework that incorporates process-aware graph rewards and stratified clipping advantage estimation to further improve reasoning quality and training stability.

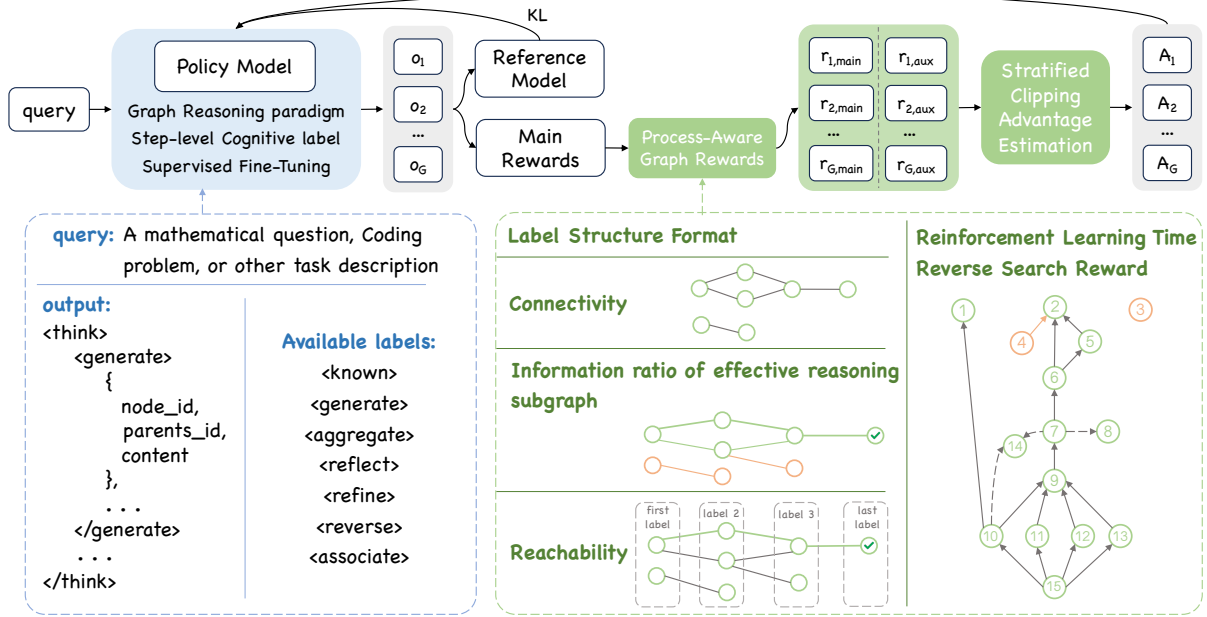


Figure 2: Overview of the proposed method.

3.1 Graph Reasoning Paradigm

We design a graph-structured reasoning paradigm in which the model directly outputs its reasoning process as a graph. This paradigm is realized through step-level cognitive labels, node-based reasoning representation, and data synthesis followed by supervised fine-tuning (SFT).

3.1.1 Step-level Cognitive Labels

Existing large reasoning models employ only coarse-grained `<think>` and `<answer>` tags, without further decomposition of the reasoning process. As a result, emergent operations such as reflection may occur at arbitrary stages and with inconsistent content during reasoning. To address this limitation, we introduce step-level reasoning tags. Details are provided in Appendix B.

By explicitly selecting a cognitive label at each step, the model is encouraged to reason about *how* to emerge, not only *can* emerge, which improves the interpretability of the reasoning process and the reliability of utilizing different reasoning patterns.

3.1.2 Node-based Reasoning Representation

Under the graph reasoning paradigm, each cognitive label wraps one or more reasoning nodes. Each node is defined as a tuple:

$$v = (id, parents_id, content), \quad (1)$$

where *id* is the unique identifier of the node, *parents_id* is a list of parent node ids, and *content*

denotes the content of the reasoning step.

This node-based representation enables the reasoning process to be abstracted into a directed graph, which forms the foundation for graph-structured, process-aware optimization.

3.1.3 Data Synthesis and SFT

To enhance the model’s structural reasoning while maintaining its inherent inference fluidity, we implement an iterative synthesis pipeline. This process consists of three primary stages:

Generating and Validating Reasoning Traces.

For each problem p in the dataset \mathcal{P} , we first sample a reasoning trace using a teacher LLM \mathcal{T} and verify its final answer against the ground truth. If the answer is correct, the trace is retained for subsequent graph-structured transformation. If the answer is incorrect, we trigger a regeneration process with an upper bound on the number of retries. This regeneration continues until either a correct reasoning trace is obtained or the maximum retry limit is reached. Only reasoning traces with verified correct answers are passed to the next stage.

Graph-Structured Translation.

Given a validated reasoning trace T , we translate it into the Graph Reasoning Paradigm. This translation explicitly exposes the latent reasoning structure while preserving the original semantic content of the trace. Graph representations derived from originally correct and regenerated traces are jointly col-

lected for quality inspection. Detailed prompts and examples are provided in Appendix C.

Graphical CoT Verification and Refinement.

To ensure the quality of the synthesized graph-structured reasoning data, we introduce an automated quality control loop that evaluates each generated graph from two complementary aspects:

- **Node–Label Consistency.** We verify that the reasoning content of each node conforms to the semantic scope of its assigned label.
- **Parent–Child Coherence.** We check whether each node is logically coherent with its parent node(s), ensuring a consistent and non-contradictory reasoning flow.

If a graph G fails either of the above checks, we generate structured feedback that explicitly identifies the problematic nodes and the corresponding violation types. This feedback, together with a pre-defined refinement prompt template, is fed back to the LLM to trigger a re-translation step. The refinement process iterates until the graph passes all quality checks or a maximum number of translation attempts is reached. Detailed prompts and cases are provided in Appendix D.

Supervised Fine-Tuning with Mixed Reasoning

Formats. After constructing the graph-structured reasoning dataset, we perform supervised fine-tuning on both graph-structured traces and a subset of original CoT traces that could not be converted into valid graphs. Although these samples do not conform to the graph format, their problem–solution pairs and coherent reasoning steps remain informative. This mixed-format training allows the model to internalize the graph reasoning paradigm while preserving fluency, completeness, and generalization of its native reasoning behavior.

3.2 PASC-GRPO

Fine-grained control over the reasoning process is crucial for solving complex reasoning tasks. (Shao et al., 2025) After SFT, the model internalizes the graph reasoning paradigm, where the reasoning process can be abstracted as a graph independent of specific reasoning content. Building on it, we propose **PASC-GRPO (Process-Aware Stratified Clipping GRPO)**, which enables optimization of reasoning length and process quality.

3.2.1 Process-Aware Graph Rewards

We use NetworkX (Hagberg et al., 2008) to construct reasoning graphs. On this basis, we design *Process-Aware Graph Rewards*, which can be grouped into three categories:

- **Graph structure format rewards**, including the *Label Structure Format Reward*, which enforce valid and well-formed reasoning graphs.
- **Reasoning length rewards**, including the *Connectivity Reward*, which encourages fewer connected components, and the *Information Ratio of Effective Reasoning Subgraph*, which promotes shorter reasoning paths within each component.
- **Reasoning process quality rewards**, including the *Reachability Reward* at the global process level, and the *Reinforcement Learning Time Reverse Search Reward* at the step level.

Label Structure Format Rewards. We introduce it to enforce the structural validity of graph-structured reasoning. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote the reasoning graph, where $Pa(v)$ is the set of parent nodes of v . Let \mathcal{K} , \mathcal{A} , and \mathcal{R} denote *known*, *aggregate*, and *refine* nodes, and \mathcal{T} denote the set of reasoning tags. The overall format reward R_{fmt} is defined as the average of three complementary sub-rewards:

$$R_{fmt} = \frac{1}{3} (R_{dens} + R_{topo} + R_{para}). \quad (2)$$

- **Node Density.** For *aggregate* and *refine*, each tag is required to wrap exactly one node:

$$R_{dens} = \frac{1}{|\mathcal{A} \cup \mathcal{R}|} \sum_{v \in \mathcal{A} \cup \mathcal{R}} \mathbb{1}(|v|_{\text{label}} = 1). \quad (3)$$

- **Topological Validity.** Specifically, *known* nodes must not depend on prior reasoning, *aggregate* nodes must combine multiple predecessors, and *refine* nodes must extend exactly one preceding step:

$$R_{topo} = \frac{1}{|\mathcal{K} \cup \mathcal{A} \cup \mathcal{R}|} \sum_{v \in \mathcal{K} \cup \mathcal{A} \cup \mathcal{R}} \Phi(v), \quad (4)$$

$$\Phi(v) = \begin{cases} 1, & |Pa(v)| = 0, v \in \mathcal{K}, \\ 1, & |Pa(v)| > 1, v \in \mathcal{A}, \\ 1, & |Pa(v)| = 1, v \in \mathcal{R}, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

- **Parallelism.** Nodes within a tag are treated as parallel reasoning units and are therefore prohibited from forming parent–child relations:

$$R_{para} = \frac{1}{|\mathcal{T}|} \sum_{T \in \mathcal{T}} \mathbb{1} \left(\forall v_i, v_j \in T, v_i \notin Pa(v_j) \right) \quad (6)$$

Connectivity Reward. To discourage fragmented reasoning, we define the connectivity reward based on the number of connected subgraphs:

$$R_{conn} = \frac{1}{n}, \quad (7)$$

where n is the number of connected subgraphs in \mathcal{G} . This reward favors reasoning with fewer isolated branches, resembling more focused reasoning.

Information Ratio of Effective Reasoning Subgraph. We define the effective reasoning subgraph (ERS) as the part of the reasoning graph that effectively contributes to reaching the final answer. The shortest path from the initial conditions to the answer node forms a backbone of effective reasoning. Reasoning branches that diverge from this backbone but eventually merge back are also considered effective. Branches that do not reconnect are treated as ineffective explorations, as they may correspond to incorrect reasoning attempts.

$$\mathcal{V}_{ERS} = \{ v \in \mathcal{V} \mid v_{start} \rightsquigarrow v \rightsquigarrow v_{end} \}. \quad (8)$$

The ERS Information Ratio R_{ers} is measured by token count $I(\cdot)$:

$$R_{ent} = \frac{\sum_{v \in \mathcal{V}_{ERS}} I(v)}{\sum_{v \in \mathcal{V}_{Total}} I(v)} \quad (9)$$

This encourages the model to avoid "dead-end" branches and redundant chatter.

Reachability Reward. To mitigate the "wrong process, correct answer" phenomenon (Akter et al., 2025), we evaluate the end-to-end logical flow:

$$R_{reach} = \mathbb{1}(v_{start} \rightsquigarrow v_{end}) \quad (10)$$

where $R_{reach} = 1$ if there exists a directed path in \mathcal{G} , ensuring the reasoning process is complete and leads to the final answer.

Reinforcement Learning Time Reverse Search Reward. Unlike forward heuristics such as MCTS or PRM, which incrementally expand reasoning paths to approximate the correct answer, our reward directly leverages the known answer. After the reasoning graph \mathcal{G} is completed, we traverse

it backward from the answer node v_{ans} to identify which nodes actually contribute to reaching the answer. For each node $v \in \mathcal{V}$, we assign

$$r(v) = \begin{cases} 1 & \text{if } v \rightsquigarrow v_{ans}, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

where $v \rightsquigarrow v_{ans}$ indicates that v is reachable from v_{ans} via backward traversal. The total reward is computed as the average over all nodes:

$$R_{rev} = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} r(v). \quad (12)$$

Overall Process-Aware Graph Reward. The final process-aware graph reward is defined as a weighted sum with scalar weights w_1, \dots, w_5 , where each component reward is normalized to $[0, 1]$ and $\sum_{i=1}^5 w_i = 1$:

$$R_{graph} = w_1 R_{fmt} + w_2 R_{conn} + w_3 R_{ers} + w_4 R_{reach} + w_5 R_{rev} \quad (13)$$

3.2.2 Stratified Clipping Advantage Estimation

To mitigate the reward hacking phenomenon — where the model might optimize graph-structured rewards despite producing incorrect answers — we propose the **Stratified Clipping Advantage Estimation** (SCAE) method. This approach hierarchically prioritizes task accuracy over auxiliary structural signals. The logical flow of the proposed SCAE is illustrated in Figure 3.

Group Stratification. For a sampled group of G reasoning traces, we first calculate the mean accuracy reward $\bar{r}_{acc} \in [0, 1]$. The group is then partitioned into a **Correct Group** (\mathcal{G}_{corr}) where $r_{acc}^{(i)} = 1$ and a **Wrong Group** (\mathcal{G}_{wrong}) where $r_{acc}^{(j)} = 0$. We define the baseline accuracy advantages for these strata as:

$$A_{acc} = \begin{cases} 1 - \bar{r}_{acc}, & \text{if } i \in \mathcal{G}_{corr} \\ 0 - \bar{r}_{acc}, & \text{if } j \in \mathcal{G}_{wrong} \end{cases} \quad (14)$$

This ensures the baseline advantage is non-negative ($A_{acc} \geq 0$) for the correct group and non-positive ($A_{acc} \leq 0$) for the wrong group.

Asymmetrical Auxiliary Clipping. Within each stratum, we calculate the mean of the auxiliary graph rewards, denoted as \bar{r}_{aux}^+ for the correct group and \bar{r}_{aux}^- for the wrong group. Here, the auxiliary reward refers to all rewards other than the main reward (e.g. accuracy reward), such as

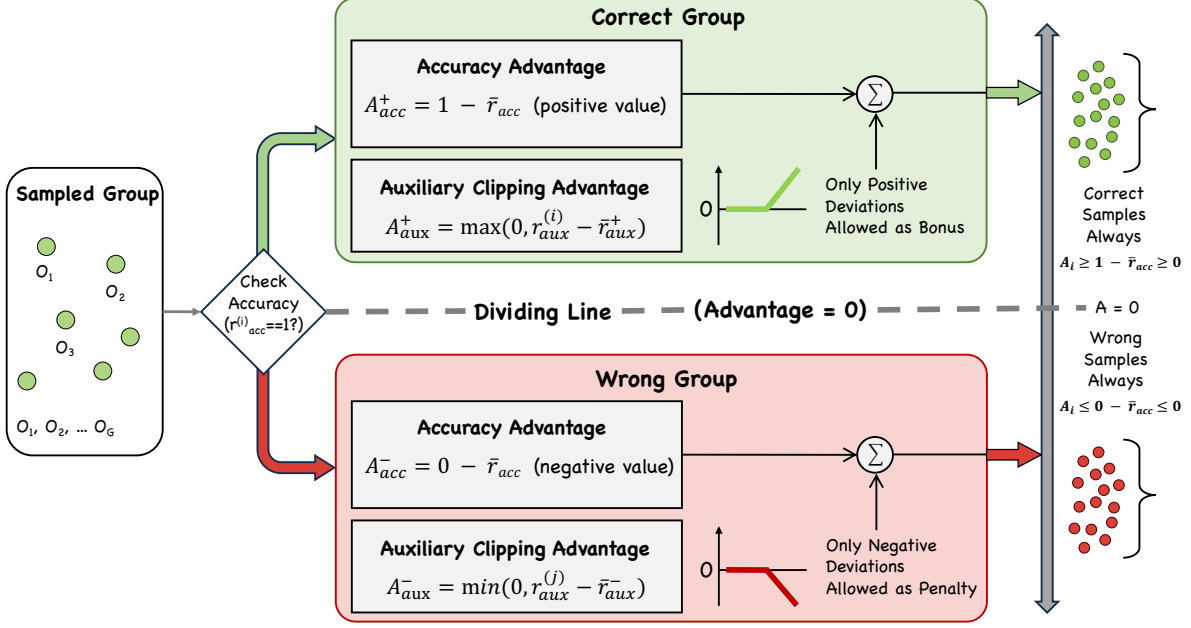


Figure 3: Illustration of the Stratified Clipping Advantage Estimation.

graph or format rewards. To prevent structural rewards from overriding the accuracy signal, we apply asymmetrical clipping:

- **Correct Group Clipping:** In \mathcal{G}_{corr} , structural rewards are treated strictly as bonuses. Even if a trace’s structural reward $r_{aux}^{(i)}$ is below the group mean, it is not penalized, ensuring the final advantage A_i never falls below the accuracy baseline:

$$A_i = (1 - \bar{r}_{acc}) + \max(0, r_{aux}^{(i)} - \bar{r}_{aux}^+) \quad (15)$$

to ensure $i \in \mathcal{G}_{corr}$, $A_i \geq 1 - \bar{r}_{acc} \geq 0$.

- **Wrong Group Clipping:** In \mathcal{G}_{wrong} , structural rewards serve exclusively as penalties. No matter how high the structural quality $r_{aux}^{(j)}$ is relative to the mean, it receives no credit, ensuring the advantage remains non-positive:

$$A_j = (0 - \bar{r}_{acc}) + \min(0, r_{aux}^{(j)} - \bar{r}_{aux}^-) \quad (16)$$

to ensure $j \in \mathcal{G}_{wrong}$, $A_j \leq 0 - \bar{r}_{acc} \leq 0$.

In summary, SCAE offers the following advantages for training graph-reasoning models:

- **Accuracy Primacy:** By setting a hard floor for correct samples ($A_i \geq A_{acc}^+ \geq 0$) and a hard ceiling for incorrect samples ($A_j \leq A_{acc}^- \leq 0$), SCAE ensures that correctness remains the dominant optimization objective.

- **Reward Hacking Resilience:** The asymmetric clipping mechanism prevents the model from "cheating" by generating high-quality graph structures for incorrect derivations to offset the accuracy penalty.

4 Experiments

4.1 Experimental Setup

Baselines. We include LLaMA-3.1-8B-Instruct from Meta (Dubey et al., 2024) and Gemma-3-12B-IT / Gemma-3-27B-IT from Google (Kamath et al., 2025) as comparison baselines. To evaluate the effectiveness of GRP and PASC-GRPO, we train on the pre-trained Qwen3-4B-Base and Qwen3-8B-Base models from Qwen (Yang et al., 2025a). Detailed training data construction procedures are provided in Appendix E.

Benchmarks. For mathematical reasoning, the benchmarks includes GSM8K (Cobbe et al., 2021), MATH500 (Lightman et al., 2023), and competition-level benchmarks such as AMC23 (Math AI Dataset Collection, 2025c), AIME24 (Math AI Dataset Collection, 2025a), and AIME25 (Math AI Dataset Collection, 2025b). For code generation, we utilize the widely-used MBPP (Austin et al., 2021), MBPP+ (Liu et al., 2023), HumanEval (Chen et al., 2021), HumanEval+ (Liu et al., 2023), and the highly challenging LiveCodeBench v5 (Jain et al., 2024).

Table 1: Experimental results on Mathematical Reasoning and Code Generation benchmarks.

Mathematical Reasoning					
Model	GSM8K	MATH500	AMC23	AIME24	AIME25
LLaMA-3.1-8B-Instruct	85.00	54.80	41.20	6.30	2.70
Gemma-3-12B-IT	94.40	85.60	77.30	22.40	18.80
Gemma-3-27B-IT	95.90	90.00	80.50	32.60	24.00
Qwen3-4B-Base	70.20	55.43	19.17	10.00	6.33
+ GRP-SFT	82.30	72.12	63.33	26.67	18.70
+ PASC-GRPO	90.45	85.32	71.43	33.90	24.13
Qwen3-8B-Base	73.01	60.08	38.12	10.00	7.60
+ GRP-SFT	87.72	81.60	75.00	40.00	33.33
+ PASC-GRPO	95.37	91.20	82.50	46.67	38.79
Code Generation					
Model	MBPP	MBPP+	HumanEval	HumanEval+	LiveCodeBench
LLaMA-3.1-8B-Instruct	61.20	52.30	69.70	62.80	10.80
Gemma-3-12B-IT	73.00	62.10	85.40	78.20	25.70
Gemma-3-27B-IT	74.40	63.00	87.80	80.00	26.90
Qwen3-4B-Base	62.40	51.30	75.60	70.70	21.14
+ GRP-SFT	65.43	54.2	81.67	76.90	42.93
+ PASC-GRPO	67.33	55.43	82.45	77.83	46.79
Qwen3-8B-Base	73.27	61.73	80.47	74.27	29.76
+ GRP-SFT	75.45	64.71	86.79	82.09	53.93
+ PASC-GRPO	77.49	65.24	88.43	83.98	56.12

Evaluation Metrics. For larger datasets (GSM8K, MATH500, and LiveCodeBench), we report the avg@3 accuracy; for other benchmarks, we report the avg@16 accuracy. Considering the extensive derivation required for competition-level mathematics (AIME24, AIME25), we set the max_completion_tokens to 32K, while maintaining a limit of 8K for other datasets.

4.2 Main Results

Performance Overview. Table 1 reports the performance of our models on mathematical reasoning and code generation benchmarks. Both GRP-SFT and PASC-GRPO consistently improve accuracy, demonstrating that the proposed GRP effectively enhances reasoning capabilities.

Impact of GRP-SFT. GRP-SFT brings substantial gains in mathematical reasoning. For Qwen3-8B-Base, accuracy improves by 21.52% on MATH500 and 18.30% on GSM8K, with a particularly large gain of 36.88% on the competition-level AMC23 benchmark. Code generation tasks also benefit from GRP-SFT, indicating that the proposed graph-based reasoning paradigm generalizes beyond mathematics.

Effect of PASC-GRPO. PASC-GRPO further enhances performance, especially on high-difficulty benchmarks. On AIME24, Qwen3-4B and Qwen3-8B achieve additional gains of 7.23% and 6.67% over the SFT stage, respectively. These results suggest that the process-aware reward are effective for multi-step problems.

Comparison with Open-Source LLMs. Despite using fewer parameters and training resources, our best model surpasses widely adopted open-source models on competition-level benchmarks. In particular, it surpasses Gemma-3-27B-IT by over 14% accuracy on both AIME24 and AIME25, and achieves comparable or better performance on standard math and code benchmarks such as MATH500, AMC23, MBPP+, and HumanEval+.

4.3 Ablation Studies

Ablation of GRP and PASC-GRPO. We study the effectiveness of GRP and PASC-GRPO. For SFT, we use the same data as GRP-SFT but replace graph-structured reasoning with standard chain-of-thought. As shown in Table 2, this variant performs

Table 2: Ablation Experiment Results. Accuracy (%).

Method	MATH500	AMC23	AIME24
Qwen3-8B-Base	60.08	38.12	10.00
GRP-SFT	81.60	75.00	40.00
SFT	76.45	67.80	31.50
PASC-GRPO	91.20	82.50	46.67
GRPO	88.75	79.85	43.80
w/o Reachability	87.45	79.30	40.15
w/o Reach. & RSR	83.12	76.55	35.80
w/o SCAE	84.15	76.20	32.50

worse than GRP-SFT on all benchmarks.

For RL, we compare PASC-GRPO with the original GRPO under identical training settings. GRPO consistently yields lower performance. These results show that both graph-structured reasoning and process-aware rewards are essential for strong reasoning performance.

Ablation of Reasoning-Length Rewards. We remove the connectivity reward and further remove ERS Ratio, starting from the GRP-SFT baseline. Figure 4 shows the changes in accuracy and length during training, comparing the cases with and without connectivity and ERS Ratio rewards. For more details on the ablation studies, see Appendix F.

Ablation of Reasoning process quality rewards. Table 2 reports the ablation results of the reachability and reverse search rewards, starting from the GRP-SFT baseline. Removing the reachability reward leads to a performance drop across all benchmarks. Further removing the reverse search reward causes a more substantial degradation. These results indicate that both rewards are critical for guiding effective graph reasoning.

Ablation of SCAE. Table 2 shows that using the original GRPO instead of Stratified Clipping Advantage Estimation (SCAE) consistently degrades performance across all benchmarks. These results suggest that SCAE is crucial for stabilizing the reasoning process and improving accuracy.

4.4 Discovery

As shown in Table 3, our best model trained from Qwen3-8B-Base matches or surpasses the performance of Qwen3-8B on hard code generation benchmarks. Qwen3-8B is a state-of-the-art model distilled using extensive human and computational resources. These results highlight the effectiveness of our methods in handling complex reasoning

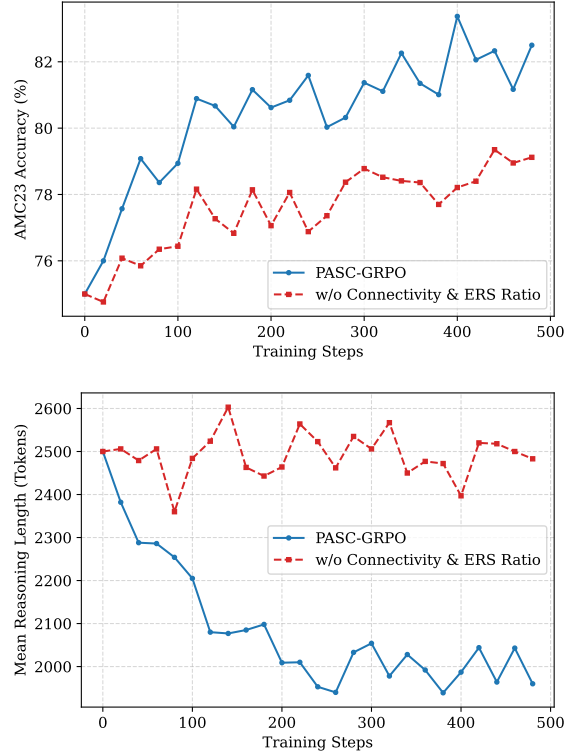


Figure 4: (Top) Accuracy comparison between PASC-GRPO and its variants without connectivity and ERS Ratio rewards on AMC23. (Bottom) Corresponding changes in mean reasoning length during training.

tasks, especially code generation.

Table 3: Performance Comparison on Code Generation

Model	MBPP	MBPP+	LiveCodeBech
qwen3-8B	76.70	65.10	57.29
qwen3-8B-Base	73.27	61.73	29.76
+ GRP-SFT	75.45	64.71	53.93
+ PASC-GRPO	77.49	65.24	56.12

5 Conclusion

In this paper, we propose the Graph Reasoning Paradigm (GRP), which shifts LLM reasoning from plain text into structured, symbolic graph representations. We further introduce PASC-GRPO, a reinforcement learning framework that leverages topology-aware rewards and Stratified Clipping Advantage Estimation to enhance reasoning quality while mitigating reward hacking. Our experiments across math and coding demonstrate that GRP significantly improves performance. This work underscores the potential of structured symbolic reasoning and topology-aware rewards in developing more reliable and efficient reasoning models.

541 Limitations

542 While the Graph Reasoning Paradigm and PASC-
543 GRPO demonstrate substantial improvements in
544 mathematical reasoning and code generation, sev-
545 eral promising directions remain for future explora-
546 tion. Currently, our evaluation primarily focuses
547 on these logic-intensive domains; extending the
548 graph-structured thinking framework to broader
549 tasks, such as commonsense reasoning or agentic
550 planning, presents a valuable opportunity to fur-
551 ther assess its generalization capabilities. Addition-
552 ally, while we have verified the effectiveness of
553 our approach on 4B and 8B parameter models, in-
554 vestigating the scaling laws of graph reasoning on
555 larger-scale foundation models could reveal further
556 performance gains. Finally, future work may also
557 explore more dynamic and adaptive graph topol-
558 ogy generation methods to support increasingly
559 complex and flexible cognitive processes without
560 relying on predefined cognitive label sets.

561 References

562 Sanjeda Akter, Ibne Farabi Shihab, and Anuj Sharma.
563 2025. [Inducing faithfulness in structured rea-
564 soning via counterfactual sensitivity](#). *ArXiv*,
565 abs/2509.01544.

566 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten
567 Bosma, Henryk Michalewski, David Dohan, Ellen
568 Jiang, Carrie J. Cai, Michael Terry, Quoc V. Le, and
569 Charles Sutton. 2021. [Program synthesis with large
570 language models](#). *ArXiv*, abs/2108.07732.

571 Maciej Besta, Nils Blach, Aleš Kubíček, Robert Ger-
572 stenberger, Lukas Gianinazzi, Joanna Gajda, Tomasz
573 Lehmann, Michal Podstawski, Hubert Niewiadom-
574 ski, Piotr Nyczyk, and Torsten Hoefler. 2023. [Graph
575 of thoughts: Solving elaborate problems with large
576 language models](#). In *AAAI Conference on Artificial
577 Intelligence*.

578 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming
579 Yuan, Henrique Pondé, Jared Kaplan, Harrison Ed-
580 wards, Yura Burda, Nicholas Joseph, Greg Brockman,
581 Alex Ray, Raul Puri, Gretchen Krueger, Michael
582 Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin,
583 Brooke Chan, Scott Gray, and 34 others. 2021. [Eval-
584 uating large language models trained on code](#). *ArXiv*,
585 abs/2107.03374.

586 Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng,
587 Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang
588 Zhou, Te Gao, and Wangxiang Che. 2025. [To-
589 wards reasoning era: A survey of long chain-of-
590 thought for reasoning large language models](#). *ArXiv*,
591 abs/2503.09567.

Karl Cobbe, Vineet Kosaraju, Mo Bavarian, Mark
592 Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plap-
593 pert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano,
594 Christopher Hesse, and John Schulman. 2021. [Train-
595 ing verifiers to solve math word problems](#). *ArXiv*,
596 abs/2110.14168. 597

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang,
598 Jun-Mei Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
599 Shirong Ma, Peiyi Wang, Xiaoling Bi, Xiaokang
600 Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou,
601 Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 179 oth-
602 ers. 2025. [Deepseek-r1: Incentivizing reasoning ca-
603 pability in llms via reinforcement learning](#). *ArXiv*,
604 abs/2501.12948. 605

Carson E. Denison, Monte Stuart MacDiarmid, Fazl
606 Barez, David Kristjanson Duvenaud, Shauna Kravec,
607 Samuel Marks, Nicholas Schiefer, Ryan Sok-
608 laski, Alex Tamkin, Jared Kaplan, Buck Shlegeris,
609 Samuel R. Bowman, Ethan Perez, and Evan Hub-
610 inger. 2024. [Sycophancy to subterfuge: Investigating
611 reward-tampering in large language models](#). *ArXiv*,
612 abs/2406.10162. 613

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey,
614 Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman,
615 Akhil Mathur, Alan Schelten, Amy Yang, Angela
616 Fan, Anirudh Goyal, Anthony S. Hartshorn, Aobo
617 Yang, Archi Mitra, Archie Sravankumar, Artem Ko-
618 renev, Arthur Hinsvark, Arun Rao, Aston Zhang, and
619 510 others. 2024. [The llama 3 herd of models](#). 620

Ahmed El-Kishky. 2024. [Openai o1 system card](#). *ArXiv*,
621 abs/2412.16720. 622

Ahmed El-Kishky, Alexander Wei, Andre Saraiva, Bo-
623 rys Minaev, Daniel Selsam, David Dohan, Francis
624 Song, Hunter Lightman, Ignasi Clavera, Jakub W. Pa-
625 chocki, Jerry Tworek, Lorenz Kuhn, Lukasz Kaiser,
626 Mark Chen, Max Schwarzer, Mostafa Rohaninejad,
627 Nat McAleese, o3 contributors, Oleg Murk, and 5
628 others. 2025. [Competitive programming with large
629 reasoning models](#). *ArXiv*, abs/2502.06807. 630

Dileep George, Rajeev Vijay Rikhye, Nishad Gothoskar,
631 J. Swaroop Guntupalli, Antoine Dedieu, and Miguel
632 Lázaro-Gredilla. 2021. [Clone-structured graph rep-
633 resentations enable flexible learning and vicarious
634 evaluation of cognitive maps](#). *Nature Communica-
635 tions*, 12. 636

Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart.
637 2008. Exploring network structure, dynamics, and
638 function using networkx. In *Proceedings of the
639 7th Python in Science Conference*, pages 11 – 15,
640 Pasadena, CA USA. 641

Tao He, Hao Li, Jingchang Chen, Runxuan Liu, Yixin
642 Cao, Lizi Liao, Zihao Zheng, Zheng Chu, Jiafeng
643 Liang, Ming Liu, and Bing Qin. 2025. [Breaking the
644 reasoning barrier a survey on llm complex reasoning
645 through the lens of self-evolution](#). In *Annual Meeting
646 of the Association for Computational Linguistics*. 647

648	Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. Live-codebench: Holistic and contamination free evaluation of large language models for code . <i>ArXiv</i> , abs/2403.07974.	Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. 2023. Cognitive architectures for language agents . <i>Trans. Mach. Learn. Res.</i> , 2024.	701 702 703 704
654	Gemma Team Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ram'e, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean-Bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, Gael Liu, and 191 others. 2025. Gemma 3 technical report . <i>ArXiv</i> , abs/2503.19786.	Yuxi Xie, Anirudh Goyal, Wenye Zheng, Min-Yen Kan, Timothy P. Lillicrap, Kenji Kawaguchi, and Michael Shieh. 2024. Monte carlo tree search boosts reasoning via iterative preference learning . <i>ArXiv</i> , abs/2405.00451.	705 706 707 708 709
662	Hunter Lightman, Vineet Kosaraju, Yura Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let's verify step by step . <i>ArXiv</i> , abs/2305.20050.	Fengli Xu, Qianye Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Chenyang Shao, Yuwei Yan, Qinglong Yang, Yiwen Song, Si-jian Ren, Xinyuan Hu, Yu Li, J. Feng, Chen Gao, and Yong Li. 2025. Towards large reasoning models: A survey of reinforced reasoning with large language models . <i>ArXiv</i> , abs/2501.09686.	710 711 712 713 714 715 716 717
667	Jiawei Liu, Chun Xia, Yuyao Wang, and Lingming Zhang. 2023. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation . <i>ArXiv</i> , abs/2305.01210.	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. Qwen3 technical report . <i>ArXiv</i> , abs/2505.09388.	718 719 720 721 722 723 724
671	Math AI Dataset Collection. 2025a. Aime 2024 dataset. https://huggingface.co/datasets/math-ai/aime24 . Accessed: 2026-01-03.	Shu Yang, Junchao Wu, Xin Chen, Yunze Xiao, Xinyi Yang, Derek Wong, and Di Wang. 2025b. Understanding aha moments: from external observations to internal mechanisms . <i>ArXiv</i> , abs/2504.02956.	725 726 727 728
674	Math AI Dataset Collection. 2025b. Aime 2025 dataset. https://huggingface.co/datasets/math-ai/aime25 . Accessed: 2026-01-03.	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models . <i>ArXiv</i> , abs/2305.10601.	729 730 731 732 733
677	Math AI Dataset Collection. 2025c. Amc 2023 dataset. https://huggingface.co/datasets/math-ai/amc23 . Accessed: 2026-01-03.	Zhaojian Yu, Yinghao Wu, Yilun Zhao, Arman Cohan, and Xiao-Ping Zhang. 2025. Z1: Efficient test-time scaling with code . <i>ArXiv</i> , abs/2504.00810.	734 735 736
680	Ivan Moshkov, Darragh Hanley, Ivan Sorokin, Shubham Toshniwal, Christof Henkel, Benedikt Schifferer, Wei Du, and Igor Gitman. 2025. Aimo-2 winning solution: Building state-of-the-art mathematical reasoning models with openmathreasoning dataset . <i>ArXiv</i> , abs/2504.16891.	Lifan Yuan, Wendi Li, Huayu Chen, Ganqu Cui, Ning Ding, Kaiyan Zhang, Bowen Zhou, Zhiyuan Liu, and Hao Peng. 2024. Free process rewards without process labels . <i>ArXiv</i> , abs/2412.01981.	737 738 739 740
686	Zhihong Shao, Yu-Wei Luo, Chengda Lu, Zehui Ren, Jiewen Hu, Tian Ye, Zhibin Gou, Shirong Ma, and Xiaokang Zhang. 2025. Deepseekmath-v2: Towards self-verifiable mathematical reasoning .	Kaiyan Zhang, Yuxin Zuo, Bingxiang He, Youbang Sun, Runze Liu, Che Jiang, Yuchen Fan, Kai Tian, Guoli Jia, Peng Li, Yu Fu, Xingtai Lv, Yuchen Zhang, Sihang Zeng, Shang Qu, Hao-Si Li, Shijie Wang, Yuru Wang, Xi-Dai Long, and 20 others. 2025. A survey of reinforcement learning for large reasoning models . <i>ArXiv</i> , abs/2509.08827.	741 742 743 744 745 746 747
690	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-Mei Song, Mingchuan Zhang, Y. K. Li, Yu Wu, and Daya Guo. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models . <i>ArXiv</i> , abs/2402.03300.	Pei Zhou, Jay Pujara, Xiang Ren, Xinyun Chen, Heng-Tze Cheng, Quoc V. Le, Ed Huai hsin Chi, Denny Zhou, Swaroop Mishra, and Huaixiu Steven Zheng. 2024. Self-discover: Large language models self-compose reasoning structures . <i>ArXiv</i> , abs/2402.03620.	748 749 750 751 752 753
695	Parshin Shojaee, Iman Mirzadeh, Keivan Alizadeh-Vahid, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. 2025. The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity . <i>ArXiv</i> , abs/2506.06941.		

A Technical Background

A.1 Group Relative Policy Optimization (GRPO)

GRPO (Shao et al., 2024) is a reinforcement learning algorithm designed for training language models with group-based advantage estimation. For each question q , a group of outputs $\{o_1, o_2, \dots, o_G\}$ are sampled from the old policy model $\pi_{\theta_{old}}$ and optimizes the current policy π_θ using:

$$\mathcal{J}_{GRPO}(\theta) = \mathbb{E}_{q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(\cdot|q)} \left(\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} \left(\min \left[r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t} \right] - \beta \text{D}_{KL}(\pi_\theta \| \pi_{\text{ref}}) \right) \right). \quad (17)$$

$$r_{i,t}(\theta) = \frac{\pi_\theta(o_{i,t} | q, o_{i,<t})}{\pi_{\theta_{old}}(o_{i,t} | q, o_{i,<t})}. \quad (18)$$

$$\hat{A}_{i,t} = \tilde{r}_i = \frac{r_i - \text{mean}(r)}{\text{std}(r)}. \quad (19)$$

A.2 Outcome Rewards and Process Rewards

In reinforcement learning for LLMs, rewards can be categorized into two types based on when they are assigned during the generation process.

Outcome Rewards evaluate the final result of a complete generation sequence. Given a state-action trajectory $(s_1, a_1, \dots, s_T, a_T)$, the outcome reward is defined as $R = r(s_T)$, where only the terminal state s_T receives a reward signal. This approach is commonly used in tasks where quality can only be judged after seeing the complete output, such as code correctness or final answer accuracy.

Process Rewards, on the contrary, provide feedback at each intermediate step of the reasoning process. The cumulative reward is computed as $R = \sum_{t=1}^T r(s_t, a_t)$, where $r(s, a)$ assigns credit to each reasoning step. Process rewards enable more fine-grained supervision and can guide the model toward correct reasoning paths even when the final answer is incorrect.

A.3 Monte Carlo Tree Search (MCTS)

Monte Carlo Tree Search (Xie et al., 2024) is a heuristic search algorithm that builds a search tree by iteratively selecting, expanding, simulating, and backpropagating rewards. In the context of LLM

Step-level Cognitive Labels

- **KNOWN:** known conditions and goals of the problem;
- **GENERATE:** generate reasoning steps that moves the thinking forward
- **AGGREGATE:** merging multiple steps to derive a conclusion;
- **REFINE:** improving an existing step without substantial progress;
- **REFLECT:** reviewing a previous reasoning step;
- **REVERSE:** reasoning backward from the goal to derive a solution strategy;
- **ASSOCIATE:** drawing analogies from similar reasoning structures

Figure 5: Step-level Cognitive Labels.

reasoning, each node represents a partial generation state s , and edges represent token or sequence actions a .

The algorithm balances exploration and exploitation using the Upper Confidence Bound (UCB) formula:

$$\text{UCB}(s, a) = Q(s, a) + c \sqrt{\frac{\ln N(s)}{N(s, a)}}. \quad (20)$$

where $Q(s, a)$ is the estimated value of taking action a in state s , $N(s)$ is the visit count of state s , $N(s, a)$ is the visit count of the state-action pair, and c is the exploration constant.

B Step-level Cognitive Labels

Figure 5 illustrates the step-level cognitive labels used in our Graph Reasoning Paradigm. Each reasoning step in the structured graph is annotated with one of these labels, which represent distinct cognitive operations such as *known*, *generate*, *aggregate*, *reflect*, *refine*, *reverse*, and *associate*. These labels enable fine-grained symbolic abstraction of the reasoning process, facilitating both interpretability and structured evaluation of model behavior.

814	C Graph-Structured Translation Prompt		
815	and Examples		
816	C.1 Graph-Structured Translation Prompt		
817	We design a unified prompt for graph-structured		
818	reasoning translation, which specifies the node		
819	schema, thinking-mode tags, and dependency re-		
820	lations. The prompt is shown in Figure 6. The		
821	illustrative examples mentioned in the prompt are		
822	presented in the following subsections. The orig-		
823	inal Chain-of-Thought traces are generated us-		
824	ing Qwen3-Max-Preview, which provides explicit		
825	thinking processes, and are subsequently trans-		
826	lated into graph-structured representations using		
827	the Qwen3-Max model.		
828	C.2 Mathematical Examples		
829	We first present two mathematical examples that		
830	are originally included in the graph-structured		
831	translation prompt. The first example is designed to		
832	explicitly illustrate the use of <i>reverse thinking</i> and		
833	<i>associative thinking</i> (Figure 7). The second exam-		
834	ple demonstrates another mathematical reasoning		
835	process and is accompanied by a schematic visual-		
836	ization of its corresponding reasoning graph, where		
837	nodes represent intermediate reasoning steps and		
838	edges encode their dependency relations (Figures 8		
839	and 9).		
840	C.3 Code Generation Examples		
841	We include three code generation examples. Each		
842	example corresponds to a representative question		
843	type of code generation task.		
844	• Code Completion (Complete). This type		
845	simulates the automatic code completion scen-		
846	ario in integrated development environments		
847	(IDEs). The model is provided with a func-		
848	tion signature and a docstring as context and		
849	is required to generate the remaining func-		
850	tion body. This task reflects the core capabil-		
851	ities of base code models such as Codex and		
852	StarCoder, primarily evaluating syntactic cor-		
853	rectness and context-aware continuation, as		
854	exemplified by benchmarks like HumanEval		
855	(see Figure 10).		
856	• Instruction Following (Instruct). Also re-		
857	ferred to as instruction-based or conversa-		
858	tional code generation, this paradigm eval-		
859	uates a model’s alignment with human in-		
860	tent. The input consists of a natural language		
861	description of a programming task, and the		
	model must generate a corresponding func-		862
	tional implementation. This setting is the dom-		863
	inant evaluation protocol for instruction-tuned		864
	models such as GPT-4 and DeepSeek-Coder-		865
	Instruct, emphasizing the translation of natu-		866
	ral language semantics into executable logic		867
	(see Figure 11).		868
	• Competitive Programming (Online Judge).		869
	This type emulates algorithmic contest envi-		870
	ronments such as ACM-style competitions		871
	or online judges (e.g., LeetCode). Unlike		872
	instruction-following tasks that generate a sin-		873
	gle function, competitive programming re-		874
	quires the model to produce a complete script		875
	that handles standard input and output streams.		876
	As a result, this paradigm is widely regarded		877
	as the most challenging form of code genera-		878
	tion, demanding advanced algorithmic reason-		879
	ing as well as robust handling of input pars-		880
	ing, output formatting, and edge cases (see		881
	Figure 12).		882
	D Graphical CoT Verification and		883
	Refinement Prompt		884
	This appendix presents the prompt for graphical		885
	Chain-of-Thought (CoT) verification and refine-		886
	ment, which is shown in Figure 13. All verifica-		887
	tion and refinement results are generated using the		888
	Qwen3-Max model.		889
	E Training Data		890
	E.1 Supervised Fine-Tuning (SFT) Data		891
	• Math Reasoning. We use the full 40.3k sam-		892
	ples from the DeepScaleR dataset.		893
	• Code Generation. We curate 12k samples		894
	from the KodCode-V1-SFT-R1 subset, bal-		895
	ancing task difficulty (<i>Easy</i> : <i>Middle</i> :		896
	<i>Hard</i> = 4 : 3 : 3) and task types		897
	(<i>Complete</i> : <i>Instruct</i> : <i>Online_Judge</i> =		898
	4.5 : 4.5 : 1).		899
	E.2 Reinforcement Learning (RL) Data		900
	• Math Reasoning. We sample 5,000 prob-		901
	lems from the DeepScaleR dataset. Problems		902
	are attempted four times by Qwen3-8B and		903
	stratified by difficulty based on success count.		904
	Problems failed in all attempts are discarded,		905
	followed by stratified sampling with a ratio of		906
	1 : 1 : 1.5 : 1.5.		907

Table 4: Ablation Study Results of Length-Reducing Rewards.

Method	MATH500		AMC23		AIME24	
	Acc.	Length	Acc.	Length	Acc.	Length
PASC-GRPO	91.20	1260	82.50	1960	46.67	7040
w/o Connectivity	91.35	1850	81.55	2140	45.10	11200
w/o Conn. & ERS Ratio	90.52	1920	79.12	2483	43.25	22500

- **Code Generation.** We directly employ the KodCode-Light-RL-10K subset.

E.3 Data Leakage Check

To prevent potential data leakage, we first collected all evaluation test sets relevant to our tasks. After selecting the training data for both supervised fine-tuning and reinforcement learning, we cross-checked each training sample against these test sets. Any overlapping instances would have been removed. Our inspection confirmed that the training sets contain no examples from the test sets, ensuring a clean separation between training and evaluation data.

F Ablation Study Details

Table 4 reports the detailed ablation results of the proposed length-reducing rewards across multiple math reasoning benchmarks. We progressively remove the connectivity reward and ERS information ratio, to examine their individual contributions. The results show that removing these components leads to a substantial increase in reasoning length and a consistent degradation in accuracy, especially on more challenging benchmarks such as AIME24, highlighting their importance in controlling reasoning efficiency without sacrificing correctness.

The Prompt for Graph-Structured Translation

You are a helpful AI Assistant that provides well-reasoned and detailed responses. I will provide you with a reasoning process generated by an LLM. I need you to translate the reasoning process into a directed graph structure. Respond in the following format: `<think> \n...\n </think> \n <answer> \n...\n </answer>`

Besides, you must comply with below conditions:

1. During the `<think>` phase you should organize the chain of thought using below tags:
 - known: known conditions and goals that can be found in the question.
 - generate: from the current reasoning state, generate one or more new reasoning steps. It represents a step forward in the process of reasoning.
 - aggregate: merge multiple steps or jointly reason over them to produce a new reasoning step.
 - reflection: go back to a previous reasoning step. Used to re-examine the correctness of a step or process.
 - refine: improve the current node. It is a refined modification of a certain node's statement, without producing a substantial step forward in the reasoning process.
 - associative thinking: comparing the current reasoning graph structure with other similar graph structures, in order to facilitate the current reasoning process. For example, when solving a math problem, recalling the solution methods used in previous similar problems.
 - reverse thinking: starting from the goal of the problem, considering possible solution paths, and filtering them with the given conditions. This builds a abstract reverse reasoning path from the goal to the conditions, from the unknown to the known. At this stage, you do not need to perform specific actions to get the answer. You just need to use reverse thinking to think about the reasoning method. The specific reasoning will be performed in the following tags.
2. At each further reasoning step you must choose one of these tags and wrap that step's output with the chosen tag. For example: `<generate>...</generate>`
3. The complete think phase must start with `<known>...</known>`, and the final inference tag must include the final result of the question.
4. The tag content inside is a series of thinking steps, organized in a node based manner with `node_id` and `parents`. You need to ensure that the thinking process is coherent and effective, and ultimately these nodes can be organized into a directed graph. The format example for each node is as follows:

```
{
  node_id: The unique identifier of a node, usually an integer, increasing from 1.
  parents: A list of parent node IDs for this node, used to establish inference dependencies. If there is no parent node, you can fill in none.
  content: The content of this step
}
```
5. For the content wrapped in different tags, there are the following formal requirements:
 - known: It wraps one or more nodes, and the parents of these nodes should all be "none".
 - generate: It wraps one or more nodes, (1) If it wraps one node, the parents of this nodes should be a single node. (2) If it wraps two or more nodes, the parents of these nodes should be a same single node.
 - aggregate: It wraps one node, and the parent of this node should be multiple nodes.
 - reflection: It wraps one node, and the parent of this node should be one or more nodes. Its `parent_ids` must include the last node of the current reasoning chain.
 - refine: It wraps one node, and the parent of this node should be the last node in the current reasoning chain.
 - associative thinking: It wraps one node, and the parent of this node should be one or more nodes.
 - reverse thinking: It wraps one node, and the parent of this node should be one or more nodes.
6. If a tag contains multiple nodes, the parents of these nodes cannot contain other nodes in the tag.
7. If a tag contains multiple nodes, the nodes should be separated by commas. Within a node, different tags do not require commas and should be separated by line breaks.
8. Coding Format Rules (CRITICAL):
 - The content inside `<answer>` must be PURE Python code. Do NOT include any XML tags (like `</think>`) inside `<answer>`.
 - NO INDENTATION for top-level definitions: The `import` statements and the `def function_name(...)` line MUST start at the very beginning of the line (column 0). Do NOT add extra spaces before `def`.
 - Self-Contained: Include all necessary imports (e.g., `from typing import List`).
9. Engineering Safety Rules:
 - No Side Effects: DO NOT modify the input arguments in-place (e.g., use `sorted(nums)` instead of `nums.sort()`).
 - Strict Signature: Use the EXACT function name and argument names provided in the prompt, even if they contain typos. Do not change the API.

Please strictly follow the above format and requirements.

Below I'll give you several examples:

Figure 6: The prompt template for “Graph-Structured Translation”. Criteria 8 and 9 are specifically added to accommodate code generation task.

The First Example of Mathematical Reasoning

Question:

Find the smallest number (n) such that given any (n) distinct numbers from $(\{1, 2, 3, \dots, 999\})$, one can choose four different numbers (a, b, c, d) such that $(a + 2b + 3c = d)$.

Response:

```

<think>
  <known>
    {
      node_id:1
      parents:none
      content:The set is {1, 2, ..., 999}
    },
    {
      node_id:2
      parents:none
      content:We are to find the smallest n such
      that any subset of size n contains distinct
      a,b,c,d with  $a + 2b + 3c = d$ 
    },
    {
      node_id:3
      parents:none
      content:a, b, c, d must be distinct elements
      from the chosen n numbers
    }
  </known>
  <reverse thinking>
    {
      node_id:4
      parents:2
      content:To find the minimal n that forces
      the equation  $a + 2b + 3c = d$  to hold among
      any n-element subset, we consider the
      maximal size of a subset that avoids this
      equation; then n is one more than that
      maximum.
    }
  </reverse thinking>

```

Continuation in the second column.

Continued from the previous column.

```

<generate>
  {
    node_id:5
    parents:4
    content:Let  $S \subseteq \{1, \dots, 999\}$  be a subset
    with no four distinct elements satisfying
     $a + 2b + 3c = d$ . We want to maximize  $|S|$ .
  }
</generate>
<associative thinking>
  {
    node_id:6
    parents:5
    content:This resembles extremal
    combinatorics problems like Schur's
    theorem or sum-free sets. In such
    problems, choosing large numbers often
    avoids linear equations because the RHS
    (d) would exceed the universe bound.
  }
</associative thinking>
<generate>
  {
    node_id:7
    parents:6
    content:If we choose only large numbers,
    say all numbers  $> T$ , then  $a + 2b + 3c \geq$ 
 $1*T + 2*T + 3*T = 6T$ , so if  $6T > 999$ ,
    then  $d = a + 2b + 3c \notin \{1, \dots, 999\}$ , hence
    cannot be in S.
  }
</generate>
.....

```

I omitted the subsequent reasoning and answer generation process. In this example, you mainly learned how to use reverse thinking and associative thinking.

Figure 7: A mathematical example illustrating reverse and associative thinking.

The Second Example of Mathematical Reasoning

Question:

Find the sum of all integer bases $b > 9$ for which $17_{\{b\}}$ is a divisor of $97_{\{b\}}$

Response:

```

<think>
  <known>
    {
      node_id:1
      parents:none
      content:b>9
    },
    {
      node_id:2
      parents:none
      content:17_{b} is a divisor of 97_{b}
    },
    {
      node_id:3
      parents:none
      content:b is an integer
    }
  </known>
  <generate>
    {
      node_id:4
      parents:2
      content:17_{b}=b+7
    },
    {
      node_id:5
      parents:2
      content:97_{b}=9*b+7
    },
  </generate>
  <aggregate>
    {
      node_id:6
      parents:2,4,5
      content: 9*b+7=k(b+7),k>0,k is an integer
    },
  </aggregate>
  <generate>
    {
      node_id:7
      parents:6
      content:b=(7-7k)/(k-9),1<k<9,k is an integer
    }
  </generate>
  <associative thinking>
    {
      node_id:8
      parents:7
      content:When dealing with this type of problem
      before, I used the enumeration method, and I can
      apply the same method here as well.
    }
  </associative thinking>
  
```

Continuation in the second column.

Continued from the previous column.

```

    }
  </associative thinking>
  <aggregate>
    {
      node_id:9
      parents:1,3,7,8
      content:Next, I should enumerate k under
      the condition that 1<k<9, k is an integer,
      and calculate when b satisfies the
      condition that b>9 and b is an integer.
    },
  </aggregate>
  <generate>
    {
      node_id:10
      parents:9
      content:1.if k=2,b=1,false.
    },
  </generate>
  <reflection>
    {
      node_id:16
      parents:9
      content:7.if k=8,b=49,true.
    }
  </reflection>
  <generate>
    {
      node_id:17
      parents:6,16
      content:But wait: Also b+7=? and 9*b+7=?
      Possibly b+7=56 and 9*b+7=448?
      448/56=8 Yes.
    }
  </generate>
  <reflection>
    {
      node_id:18
      parents:10,11,12,13,14,15,16
      content:Sum=21+49=70
    }
  </reflection>
  </aggregate>
  </think>
  <answer>
    70
  </answer>
  
```

Figure 8: A mathematical example in graph-structured format.

Question:

Find the sum of all integer bases $b > 9$ for which 17_b is a divisor of 97_b

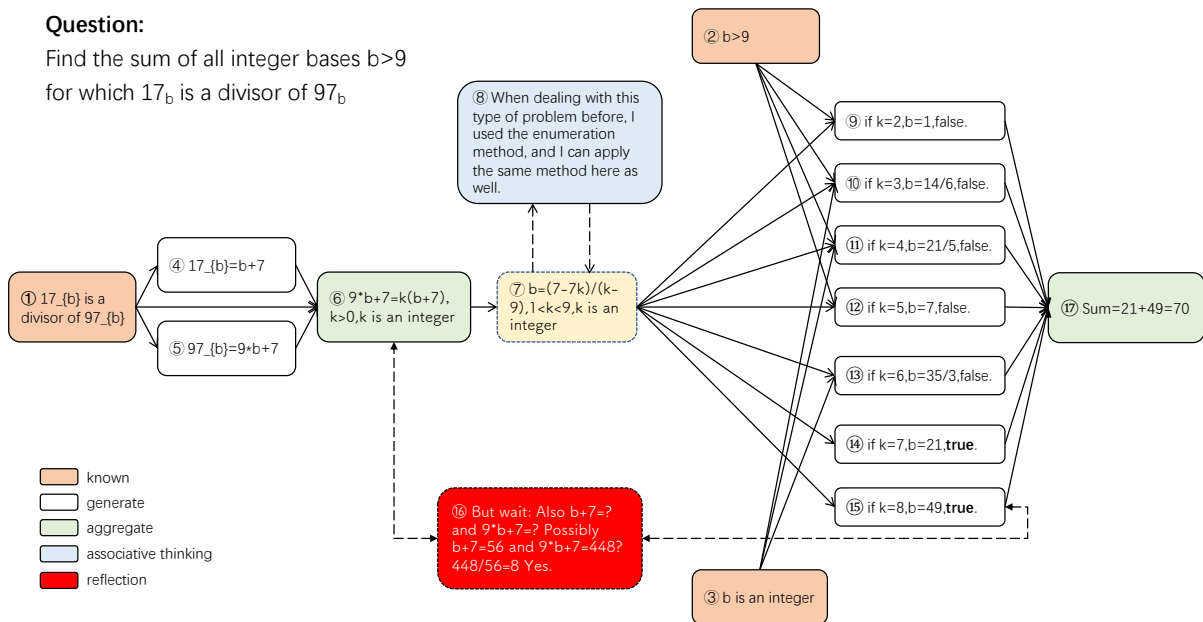


Figure 9: Graph visualization of the mathematical example.

Question:

```
def coin_change(coins: List[int], amount: int) -> int:
    """ You are given an integer array coins representing coins of different denominations
    and an integer amount representing a total amount of money. Return the fewest number
    of coins that you need to make up that amount. If that amount of money cannot be made
    up by any combination of the coins, return -1.
    >>> coin_change([1, 2, 5], 11) == 3
    >>> coin_change([2], 3) == -1
    >>> coin_change([1], 0) == 0
    >>> coin_change([1], 1) == 1
    >>> coin_change([1], 2) == 2
    """
```

Response:

```
<think>
<known>
{
  node_id: 1,
  parents: none,
  content: "Input: 'coins' (List[int]), 'amount' (int). Goal: Find minimum
  number of coins to sum to 'amount'."
},
{
  node_id: 2,
  parents: none,
  content: "Constraints: Return -1 if impossible. 'amount' >= 0 (implied
  by context/doctests). Coins are positive integers."
},
{
  node_id: 3,
  parents: none,
  content: "Edge Cases: amount is 0 (expect 0). coins list is empty
  (expect -1 unless amount is 0). amount cannot be formed."
}
</known>
<associative thinking>
{
  node_id: 4,
  parents: 1,
  content: "Pattern Recognition: This is the classic 'Coin Change'
  problem, which is a variation of the Unbounded Knapsack Problem.
  Potential approaches: BFS (shortest path in a graph) or Dynamic
  Programming (optimal substructure)."
}
</associative thinking>
<reflection>
{
  node_id: 5,
  parents: 4,
  content: "Check Greedy Approach: If we always take the largest coin
  <= amount, does it work? \nCounter-example: coins=[1, 3, 4],
  amount=6. \nGreedy: 4 + 1 + 1 = 3 coins. \nOptimal: 3 + 3 = 2 coins.
  \nConclusion: Greedy is not optimal. Must use DP or BFS."
}
</reflection>
<reverse thinking>
{
  node_id: 6,
  parents: 5,
  content: "Logic Derivation: To find dp[x] (min coins for amount x),
  the last coin added must be some coin 'c'. Thus, dp[x] = min(dp[x-c]
  + 1 for all c in coins. We solve from 0 up to amount."
}
</reverse thinking>
<generate>
{
  node_id: 7,
  parents: 6,
  content: "Plan (DP Bottom-Up):\n1. Initialize array 'dp' of size
  amount + 1.\n2. Fill 'dp' with a value > amount (representing infinity).
  dp[0] = 0.\n3. Iterate i from 1 to amount.\n4. For each coin, if i >=
  coin, dp[i] = min(dp[i], dp[i-coin] + 1).\n5. Final result is dp[amount]
  (or -1 if it's still infinity)."
}
</generate>
<refine>
{
  node_id: 8,
  parents: 7,
```

Continuation in the second column.

Continued from the previous column.

```
content: "Optimization: We can check 'coin <= i' inside the inner loop to
avoid negative index access. The 'infinity' value can be 'amount + 1' since
the worst case is all 1s (count = amount)."
```

```
</refine>
<generate>
{
  node_id: 9,
  parents: 8,
  content: "Implementation Draft:\nndef coin_change(coins, amount):\n
  dp = [amount + 1] * (amount + 1)\n  dp[0] = 0\n  for i in range(1, amount
  + 1):\n    for c in coins:\n      if i >= c:\n        dp[i] = min(dp[i], dp[i-c] + 1)\n
  return dp[amount] if dp[amount] <= amount else -1"
```

```
</generate>
<reflection>
{
  node_id: 10,
  parents: 9,
  content: "Dry Run Trace:\nInput: coins=[1, 2], amount=3.\nInit: dp=[0,
  4, 4, 4] (inf=4)\ni=1: c=1 -> dp[1]=min(4, dp[0]+1)=1. dp=[0, 1, 4,
  4]\ni=2: c=1 -> dp[2]=min(4, dp[1]+1)=2. c=2 -> dp[2]=min(2,
  dp[0]+1)=1. dp=[0, 1, 1, 4]\ni=3: c=1 -> dp[3]=min(4, dp[2]+1)=2. c=2
  -> dp[3]=min(2, dp[1]+1)=2. dp=[0, 1, 1, 2]\nResult: 2. Correct."
```

```
</reflection>
<aggregate>
{
  node_id: 11,
  parents: 2, 3, 10,
  content: "Final Verification: Handles edge case amount=0 correctly
  (dp[0]=0). Handles impossible case correctly (returns -1). Code structure
  is compliant with constraints."
}
</aggregate>
</think>
<answer>
from typing import List

def coin_change(coins: List[int], amount: int) -> int:
    """ You are given an integer array coins representing coins of different
    denominations and an integer amount representing a total amount of
    money. Return the fewest number of coins that you need to make up that
    amount. If that amount of money cannot be made up by any combination
    of the coins, return -1.
    >>> coin_change([1, 2, 5], 11) == 3
    >>> coin_change([2], 3) == -1
    >>> coin_change([1], 0) == 0
    >>> coin_change([1], 1) == 1
    >>> coin_change([1], 2) == 2
    """
    if amount < 0:
        return -1
    if amount == 0:
        return 0
    max_val = amount + 1
    dp = [max_val] * (amount + 1)
    dp[0] = 0
    for i in range(1, amount + 1):
        for coin in coins:
            if i >= coin:
                dp[i] = min(dp[i], dp[i - coin] + 1)
    return dp[amount] if dp[amount] <= amount else -1
</answer>
```

Figure 10: A code completion example.

Question:

You are given an integer array coins representing coins of different denominations and an integer amount representing a total amount of money. Return the fewest number of coins that you need to make up that amount. If that amount of money cannot be made up by any combination of the coins, return -1.

Function Signature: def coin_change(coins: List[int], amount: int) -> int:

Response:

```
<think>
<known>
{
  node_id:1,
  parents:none,
  content:Input: List of integers 'coins', Integer 'amount'. Goal:
  Find Minimum count of coins to sum to 'amount'.
},
{
  node_id:2,
  parents:none,
  content:Constraints: 'amount' can be 0 (return 0). If
  unsolvable, return -1. Coins are positive integers.
},
{
  node_id:3,
  parents:none,
  content:Edge Cases: Coins list is empty; Amount is negative;
  Large amount vs small coins (performance).
}
</known>
<reverse thinking>
{
  node_id:4,
  parents:1,
  content:To reach the target 'amount' with minimum coins, the
  last coin added must be some coin 'c' from the list. This
  means the problem reduces to finding min coins for 'amount
  - c'. This recursive structure implies we solve from 0 up to
  'amount'.
}
</reverse thinking>
<associative thinking>
{
  node_id:5,
  parents:4,
  content:This structure looks like the 'Unbounded Knapsack
  Problem'. A common approach is Greedy (always take largest)
  or Dynamic Programming. Let's try to verify if Greedy
  works because it is faster.
}
</associative thinking>
<reflection>
{
  node_id:6,
  parents:5,
  content:Verifying Greedy: Counter-example check. Coins=[1,
  3, 4], Amount=6. Greedy path: Take 4, remain 2. Take 1,
  remain 1. Take 1. Total 3 coins (4+1+1). Optimal path: Take 3,
  remain 3. Take 3. Total 2 coins (3+3). So Greedy FAILS.
  Must use Dynamic Programming.
}
</reflection>
<generate>
{
  node_id:7,
  parents:6,
  content:Plan (DP approach):
  1. Create array dp[amount + 1] initialized to 'infinity'
  (amount + 1).
  2. Set dp[0] = 0.
  3. Iterate i from 1 to amount.
```

Continuation in the second column.

Continued from the previous column.

```
4. For each coin, if i >= coin, dp[i] = min(dp[i], dp[i-coin] + 1).
}
</generate>
<aggregate>
{
  node_id:8,
  parents:2,3,7,
  content:Combine DP Plan with Edge Cases:
  If amount is 0, return 0 immediately.
  After DP loop, if dp[amount] > amount, implies no solution,
  return -1.
}
</aggregate>
<generate>
{
  node_id:9,
  parents:8,
  content:Implementation Code:
  def coin_change(coins, amount):
  if amount == 0: return 0
  max_val = amount + 1
  dp = [max_val] * (max_val)
  dp[0] = 0
  for i in range(1, amount + 1):
  for coin in coins:
  if i - coin >= 0:
  dp[i] = min(dp[i], dp[i - coin] + 1)
  return dp[amount] if dp[amount] != max_val else -1
}
</generate>
<reflection>
{
  node_id:10,
  parents:9,
  content:Dry Run Trace: Input: coins=[1, 2], amount=3. Init: dp=[0,
  4, 4, 4]. i=1: coin=1 -> dp[1]=min(4, dp[0]+1)=1. dp=[0,1,4,4].
  i=2: coin=1 -> dp[2]=min(4, dp[1]+1)=2. coin=2 -> dp[2]=min(2,
  dp[0]+1)=1. dp=[0,1,1,4]. i=3: coin=1 -> dp[3]=min(4,
  dp[2]+1)=2. coin=2 -> dp[3]=min(2, dp[1]+1)=2. dp=[0,1,1,2].
  Result: 2. Logic holds.
}
</reflection>
</think>
<answer>
from typing import List
def coin_change(coins: List[int], amount: int) -> int:
    """
    Computes the fewest number of coins needed to make up the
    amount.
    Returns -1 if impossible.
    """
    if amount == 0:
        return 0
    # Initialize DP array. amount + 1 acts as infinity.
    max_val = amount + 1
    dp = [max_val] * (amount + 1)
    dp[0] = 0
    for i in range(1, amount + 1):
        for coin in coins:
            if coin <= i:
                dp[i] = min(dp[i], dp[i - coin] + 1)
    return dp[amount] if dp[amount] <= amount else -1
</answer>
```

Figure 11: An instruction-following code generation example.

The Example of Code generation (Competitive Programming)

Question:

You are given an integer amount representing a total amount of money and a list of coins of different denominations. Write a program to calculate the fewest number of coins that you need to make up that amount. If that amount of money cannot be made up by any combination of the coins, print -1.

Input Format:

- The first line contains an integer 'amount'.
- The second line contains a sequence of integers 'coins' separated by spaces.

Output Format:

- Print a single integer representing the minimum number of coins (or -1).

Example 1: Input: 11 1 2 5 Output: 3

Example 2: Input: 3 2 Output: -1

Format: Generate an executable Python function to solve the given problem. The function should read input from 'stdin' and write the output to 'stdout'. Do not directly test on the sample inputs. Ensure that when the Python program runs, it reads the inputs, executes the algorithm, and writes the output to 'stdout'.

Response:

```
<think>
<known>
{
  node_id: 1,
  parents: none,
  content: "Task: Solve Coin Change problem (Min coins for 'amount').\nInput Source: stdin (Standard Input).\nOutput Target: stdout (Standard Output).\nInput Format: Line 1 = amount (int), Line 2 = coins (space-separated ints).",
},
{
  node_id: 2,
  parents: none,
  content: "Constraints: Print -1 if impossible. Standard integer limits. Time complexity matters (Standard OJ limits usually ~1-2s).",
}
</known>
<associative thinking>
{
  node_id: 3,
  parents: 1,
  content: "Algorithm Selection: This is the 'Coin Change' problem (Unbounded Knapsack). DP (Dynamic Programming) is the standard optimal solution. BFS is also possible but DP is easier to implement for just counting."
}
</associative thinking>
<reverse thinking>
{
  node_id: 4,
  parents: 1,
  content: "IO Strategy: Unlike function completion, we must parse raw strings.\nReading: `sys.stdin.read().split()` is robust for handling whitespace/newlines.\nParsing sequence: First token is 'amount', subsequent tokens are 'coins'."
}
</reverse thinking>
<generate>
{
  node_id: 5,
  parents: 3, 4,
  content: "Plan:\n1. Import sys.\n2. Read all stdin content.\n3. Parse 'amount' and 'coins' list.\n4. Initialize DP array 'dp' size [amount + 1] with 'amount + 1' (infinity).\n5. Set dp[0] = 0.\n6. Iterate i from 1 to amount: for each coin, update dp[i].\n7. Print result."
}
</generate>
<refine>
{
  node_id: 6,
  parents: 5,
  content: "Robustness Check: What if the coins line is empty? (Though problem implies valid coins). What if amount is 0? The DP initialization covers amount=0 correctly. We should ensure we strip inputs properly."
}
</refine>
```

Continued from the previous column.

```
<generate>
{
  node_id: 7,
  parents: 6,
  content: "Implementation Logic:\n\nimport sys\n\ndef solve():\n  input_data = sys.stdin.read().split()\n  if not input_data: return\n  amount = int(input_data[0])\n  coins = [int(x) for x in input_data[1:]]\n  # DP Logic\n  dp = [amount + 1] * (amount + 1)\n  dp[0] = 0\n  for i in range(1, amount + 1):\n    for c in coins:\n      if i >= c:\n        dp[i] = min(dp[i], dp[i-c] + 1)\n    result = dp[amount] if dp[amount] <= amount else -1\n  print(result)\n\nif __name__ == '__main__':\n  solve()
}
</generate>
<reflection>
{
  node_id: 8,
  parents: 7,
  content: "Dry Run (Mental Sandbox):\nInput:\n11\n1 2 5\nTokens: ['11', '1', '2', '5']\namount = 11\ncoins = [1, 2, 5]\nDP Init: dp size 12.\ni=1..11 loop runs.\nFinal dp[11] will be 3 (5+5+1 or similar combo).\nOutput: 3.\nInput:\n3\n2\nTokens: ['3', '2']\namount = 3, coins = [2]\ndp[1] -> inf\ndp[2] -> 1\ndp[3] -> inf (since 3-2=1, dp[1] is inf)\nOutput: -1.\nLogic holds."
}
</reflection>
<aggregate>
{
  node_id: 9,
  parents: 7, 8,
  content: "Final Polish: Combine logic into the required executable format. Ensure strict 'sys.stdin' usage and no extra prompt text (like 'Enter number:') as it breaks OJ graders."
}
</aggregate>
</think>
<answer>
import sys
def solve():
  input_data = sys.stdin.read().split()
  if not input_data:
    return
  amount = int(input_data[0])
  coins = [int(x) for x in input_data[1:]]
  # dp[i] = min coins to make i
  max_val = amount + 1
  dp = [max_val] * (amount + 1)
  dp[0] = 0
  for i in range(1, amount + 1):
    for coin in coins:
      if coin <= i:
        dp[i] = min(dp[i], dp[i - coin] + 1)
  print(dp[amount] if dp[amount] <= amount else -1)

if __name__ == "__main__":
  solve()
</answer>
```

Continuation in the second column.

Figure 12: A competitive programming code generation example.

The Prompt for Graphical CoT Verification and Refinement

You are a strict evaluator for graph-structured reasoning.

Each reasoning node is associated with:

- a label indicating its thinking mode,
- textual content,
- and optional parent node(s).

Your task is to evaluate the reasoning graph from two dimensions.

Dimension 1: Label–Content Consistency

For each node, judge whether its content complies with the definition of its assigned label.

The label definitions are as follows:

- known: Known conditions explicitly stated in the question. This node should not depend on any parent.
- generate: Generate one or more new reasoning steps from the current reasoning state.
- aggregate: Combine multiple reasoning steps to produce a new conclusion. This node should have multiple parents.
- feedback / reflection: Go back to a previous reasoning step to re-examine its correctness.
- refine: Improve or rephrase an existing reasoning step without introducing substantial new progress.
- associative thinking: Draw analogies to similar problems or reasoning graph structures to facilitate reasoning.
- reverse thinking: Start from the target goal and reason backward, filtering possible solution paths using given conditions.

Score each node with a value between 0.0 and 1.0, where:

- 1.0 indicates perfect consistency with the label definition,
- 0.0 indicates complete mismatch or contradiction.

Dimension 2: Parent–Child Semantic Coherence

For nodes with parent(s), judge whether the child node logically follows from and builds upon its parent node(s).

Evaluation criteria include:

- The child should not contradict the parent reasoning.
- Repeating or refining parent reasoning in a consistent way should receive a high score.
- Introducing a valid new reasoning step that logically follows should receive a high score.
- Contradictions or violations of parent logic should receive a low score.

If a node has no parent (e.g., a known node), it is considered coherent by default.

Score semantic coherence with a value between 0.0 and 1.0.

Output Requirement

For each evaluation, respond only with a numeric score between 0.0 and 1.0.

Figure 13: Prompt for “Graphical CoT Verification and Refinement”.