

# EXPOSING WEAKNESSES OF LARGE REASONING MODELS THROUGH GRAPH ALGORITHM PROBLEMS

000  
001  
002  
003  
004  
005 **Anonymous authors**  
006 Paper under double-blind review  
007  
008  
009  
010

## ABSTRACT

011 Large Reasoning Models (LRMs) have advanced rapidly, yet existing benchmarks  
012 on mathematics, code, and common-sense reasoning remain limited: they lack  
013 long-context evaluation, offer insufficient challenge, and provide answers that are  
014 difficult to verify programmatically. We introduce **GRALGOBENCH**, a benchmark  
015 designed to evaluate LRMs through graph algorithm problems. Such problems are  
016 particularly well-suited for probing reasoning abilities: they demand long-context  
017 reasoning, allow fine-grained control of difficulty levels, and enable standardized  
018 programmatic evaluation. Across nine tasks, our systematic experiments reveal  
019 two major weaknesses of current LRMs. **First**, accuracy deteriorates sharply with  
020 longer context **inputs**—falling below 50% once graphs exceed 120 nodes—driven  
021 by frequent execution errors, weak memory, and redundant reasoning. **Second**,  
022 LRMs suffer from an *over-thinking* phenomenon, primarily driven by extensive  
023 yet largely ineffective self-verification, which inflates reasoning traces without  
024 improving correctness. By exposing these limitations, **GRALGOBENCH** establishes  
025 graph algorithm problems as a rigorous, multidimensional, and practically relevant  
026 testbed for advancing the study of reasoning in LRMs. Code is available at <https://anonymous.4open.science/r/GrAlgoBench-7D17>.  
027  
028

## 1 INTRODUCTION

029 Large language models (LLMs) have evolved rapidly over the past few years, becoming powerful  
030 general-purpose AI tools with remarkable achievements in natural language processing and complex  
031 reasoning. Recently, the emergence of Large Reasoning Models (LRMs), such as OpenAI-O1 (Ope-  
032 nAI, 2025) and DeepSeek-R1 (Guo et al., 2025), has further pushed the frontier of LLM development.  
033 By leveraging long chains of thought enriched with human-like cognitive strategies—including  
034 self-verification, strategy-shift, and backtracking—these models show significant improvements on  
035 challenging reasoning tasks (Qin et al., 2024; Min et al., 2024; Feng et al., 2025; Dong et al., 2025).  
036 Accordingly, evaluating the reasoning capabilities of LRMs and probing their limitations, such as  
037 *over-thinking* and *under-thinking*, has become a timely research focus (Zheng et al., 2025; Chen et al.,  
038 2024c; Sun et al., 2025; Yang et al., 2025b;c).  
039

040 Current benchmarks for LRMs, while valuable, exhibit several critical deficiencies when centered  
041 on domains like *mathematical reasoning* (Petrov et al., 2025; He et al., 2024; Huang et al., 2024;  
042 Gulati et al., 2024; Glazer et al., 2024; Li et al., 2025b; Huang et al., 2025), *code generation* (Li  
043 et al., 2022; Dai et al., 2024; Yang et al., 2025b; Yu et al., 2024; Wei et al., 2025; He et al., 2025b;  
044 Zheng et al., 2025), and *common-sense reasoning* (White et al., 2024; Suzgun et al., 2022; Lin  
045 et al., 2025). First, they **lack long-context input evaluation**: existing benchmarks predominantly  
046 use short problem texts, which cannot be easily scaled to assess LRMs’ reasoning capabilities over  
047 extended contexts. As modern LRMs increasingly support longer inputs, rigorous evaluation of  
048 their performance on lengthy, complex reasoning tasks becomes essential. Second, they exhibit  
049 **inadequate challenge levels**: Current benchmarks are no longer sufficiently difficult for LRMs (e.g.,  
050 GPT-5 achieving 94.3% on AIME-2025), yet creating more challenging variants requires non-trivial  
051 human effort in problem redesign. Third, their answers are **not standardized and therefore hard**  
052 to verify programmatically: for example, in mathematical problems, solutions often incur diverse  
053 latex expressions (e.g.,  $\$$  vs.  $\backslash[$  vs.  $\backslashbegin\{align\}$ ) and admit multiple equivalent math formulations  
(e.g.,  $\frac{1}{3}$  vs.  $1/3$ , 10% vs. 0.1), which greatly increases verification effort.

054 Table 1: Comparison of our work with prior benchmarks on graph algorithm problems.  
055

056 <b>Work</b>	057 <b>Graph Size (by nodes)</b>	058 <b>Reasoning Taxonomy</b>	059 <b>LRMs' Evaluation</b>	060 <b>Real-world Data</b>
NLGraph (Wang et al., 2023)	5-35	$\times$	$\times$	$\times$
GPT4Graph (Guo et al., 2023)	10-20	$\times$	$\times$	$\times$
GraphQA (Fatem et al., 2023)	5-20	$\times$	$\times$	$\times$
LLM4DyG (Zhang et al., 2024)	5-20	$\times$	$\times$	$\times$
GraphInstruct (Luo et al., 2024)	5-35	Node, edge, graph level	$\times$	$\times$
GraphArena (Tang et al., 2024)	4-50	P and NP	$\times$	$\checkmark$
VisionGraph (Li et al., 2024)	5-35	$\times$	$\times$	$\times$
GraCore (Yuan et al., 2024)	8-30	Graph reasoning and graph understanding	$\times$	$\checkmark$
GraphOmni (Xu et al., 2025)	5-30	Node, edge, graph level	$\times$	$\times$
<b>Ours</b>	8-160	Enumeration, Exploration, Intuition	$\checkmark$	$\checkmark$

064 These collective limitations motivate a pivotal research question: *Can we gather a family of reasoning*  
065 *tasks that address the concerns above to form a more rigorous benchmarking suite for LRM*s?

066 In this work, we posit that *graph algorithm problems* (Wang et al., 2023; Chen et al., 2024a; Luo  
067 et al., 2024; Zhang et al., 2025) represent a superior and highly promising choice. They furnish  
068 several distinctive advantages: (1) **Effective Long-Context Input Reasoning**: Describing a graph  
069 typically requires listing nodes and edges, producing lengthy inputs whose solutions cannot be  
070 directly extracted but must be derived through multi-step reasoning. These properties make graph  
071 algorithm problems naturally suited for evaluating long-context reasoning. For example, OpenAI  
072 employed the *GraphWalks* (OpenAI, 2025) dataset to benchmark the long-context capabilities of  
073 GPT-5. (2) **Scalable Difficulty Level Control**: Task complexity can be precisely modulated by  
074 scaling graph parameters like node size, often leading to quadratic or even exponential growth in  
075 difficulty. As evidenced in Table 3, model accuracy deteriorates sharply as graph size surpasses 120  
076 nodes. (3) **Standardized and Programmatic Evaluation**: Outputs typically consist of integers or  
077 explicit graph elements (e.g., nodes, edges, or paths) that has no alternative representations, which  
078 permits standardized and programmatic verification.

079 Furthermore, graph algorithm problems are both highly extensible and resistant to data contamina-  
080 tion, as minor structural perturbations can generate a vast number of new, valid instances at scale.  
081 Crucially, graph algorithms are foundational to numerous real-world applications, including social  
082 networks (Ceccarello et al., 2024; Oettershagen et al., 2024), transportation systems (Ahmadian et al.,  
083 2024), and web mining (Chen et al., 2024b; Miyauchi et al., 2024; Pang et al., 2024). This practical  
084 relevance ensures that such benchmarks are not only rigorous but also grounded by real-world utility.

085 While several benchmarks have been proposed for evaluating LLMs on graph algorithm prob-  
086 lems (Wang et al., 2023; Zhang et al., 2024; Tang et al., 2024; Yuan et al., 2024), they suffer from  
087 three critical limitations (see Table 1). First, prior works focus mainly on non-reasoning LLMs,  
088 leaving the behaviors of O1-like LRM $s$  largely unexplored. Second, their graph sizes are capped  
089 at 50 nodes, leaving the advantage of scalable difficulty level control unexploited. Third, their task  
090 categorizations are often based on ad-hoc criteria such as local vs. global, graph understanding vs.  
091 graph reasoning, or P vs. NP, rather than grounded in algorithmic design principles that better capture  
092 distinct and practically relevant reasoning paradigms.

093 To address these gaps, we propose **GRALGOBENCH**, a new benchmark for evaluating LRM $s$  on  
094 graph algorithm problems. Figure 1 presents an overview of our benchmark, which consists of:

095 **① Systematic Dataset Construction:** We categorize problems into three categories—*Enumeration*,  
096 *Exploration*, and *Intuition*—each containing subproblems with different difficulty levels to probe  
097 distinct reasoning capacities. Our dataset comprises 2,700 graphs sampled directly from diverse  
098 real-world networks, ranging from 8 to 160 nodes across six scales, which increases diversity and  
099 mitigates data contamination risks.

100 **② Comprehensive Evaluation:** We evaluate not only non-reasoning models but also reasoning  
101 models (e.g., O1-like LRM $s$ ). Multiple metrics are adopted, including Pass@k, Cons@k, Z-score,  
102 and outcome efficiency, enabling multi-faceted analysis.

103 **③ Novel Research Questions:** We systematically investigate: (1) Do LRM $s$  excel at long-context  
104 reasoning (Section 3.1)? (2) What leads to over-thinking of LRM $s$  when solving graph algorithm  
105 problems (Section 3.2)?

106 Our experiments reveal two key findings. **First**, by evaluating LRM $s$  on graphs with varying node  
107 sizes as well as on fixed graphs with increasingly verbose textual descriptions, we observe a consistent

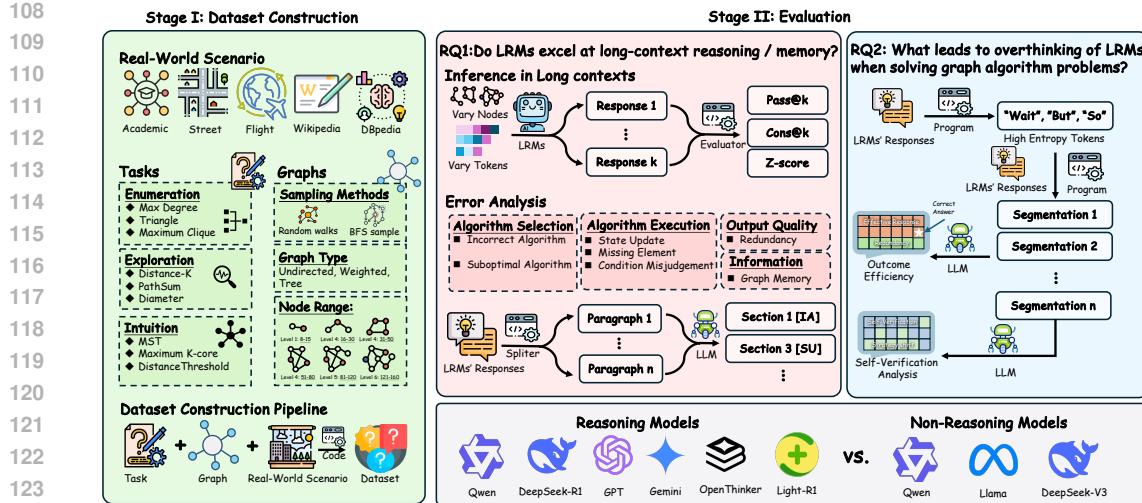


Figure 1: Benchmark overview.

127 performance decline as context length grows—for instance, most models achieve less than 50%  
128 accuracy once the graph size surpasses 120 nodes. This demonstrates that LRMAs are particularly  
129 weak at long-context reasoning. Through detailed error analysis, we identify three primary causes:  
130 frequent step-by-step execution errors, poor memory capability, and redundant reasoning patterns.  
131 **Second**, by analyzing how efficiently models reach the correct answer together with the frequency  
132 and effectiveness of self-verification, we find that LRMAs engage in extensive yet largely ineffective  
133 self-verification, which emerges as the main driver of *over-thinking*.

## 2 BENCHMARK CONSTRUCTION

137 Figure 1 illustrates the composition and structure of our dataset. Section 2.1 introduces the reasoning  
138 taxonomy, followed by task design in Section 2.2 and graph collection in Section 2.3.

### 2.1 GRAPH REASONING TAXONOMY

142 Graph algorithm problems encompass diverse reasoning paradigms, providing a comprehensive  
143 framework to evaluate LRMAs’ reasoning capabilities across multiple dimensions. Unlike prior  
144 approaches that categorize problems based on difficulty levels or by local versus global scopes (Luo  
145 et al., 2024; Tang et al., 2024; Yuan et al., 2024; Xu et al., 2025), we propose a taxonomy grounded  
146 in the algorithmic nature of reasoning, closely aligned with the major algorithmic design families in  
147 *Introduction to Algorithms (CLRS)* (Cormen et al., 2022). Specifically, graph algorithm problems  
148 can be systematically divided into three fundamental classes: Enumeration, Exploration, and  
149 Intuition. Each class corresponds directly to a canonical algorithmic design paradigm in CLRS:

- **Enumeration (brute-force algorithms):** Problems in this category align with the brute-force paradigm in CLRS. They require LRMAs to systematically generate all possible solutions or elements within a set. The canonical example in graph algorithms is brute-force enumeration of subgraphs or paths. Beyond graph problems, enumeration manifests in domains like mathematics, where LRMAs need to enumerate all permutations in a combinatorial problem or all cases in a probability calculation.
- **Exploration (search algorithms):** This category corresponds to the search paradigm in CLRS, requiring traversal of a state space with potential backtracking. LRMAs must explore multiple paths and recover from dead-ends. Typical graph examples include depth-first search and breadth-first search. Beyond graphs, it appears in tasks such as Sudoku or N-Queens, where candidate moves are tried, conflicts detected, and backtracking applied.
- **Intuition (greedy algorithms):** This category corresponds to the greedy paradigm in CLRS, where locally optimal choices are made to efficiently approximate or eventually reach a

162 global optimum. It tests LRM<sub>s</sub>’ ability to exploit subtle problem-specific signals for efficient  
 163 decisions. Representative graph examples include Dijkstra’s and Kruskal’s algorithms;  
 164 outside graphs, a classic case is Huffman coding.  
 165

166 Importantly, all three paradigms collectively test complementary reasoning abilities of LRM<sub>s</sub>: enu-  
 167 meration probes *systematic coverage of a search space*, exploration evaluates *multi-path search and*  
 168 *backtracking*, and intuition assesses *intuitive and efficient decision making*. Beyond these, all graph  
 169 problems inherently test **memory** (remembering the graph structure), **logical reasoning** (deriving  
 170 consequences from premises, e.g., inferring the existence of cycles), and **structural reasoning** (rea-  
 171 soning about relationships among nodes and edges). Owing to this diversity of reasoning paradigms  
 172 and general cognitive requirements, graph algorithm problems constitute a uniquely challenging and  
 173 informative benchmark for evaluating the reasoning capabilities of LRM<sub>s</sub>.  
 174

## 2.2 TASK DESIGN

175 With the problem taxonomy established, the next challenge lies in **classifying graph algorithm**  
 176 **problems into these categories**. It is important to emphasize that classification cannot be based  
 177 solely on a single optimal algorithm for solving the problem, as a given problem can be approached  
 178 using multiple algorithmic strategies. Moreover, we do not know which algorithmic approach an  
 179 LRM is more likely to employ when tackling a specific problem. For example, shortest path problems  
 180 can be solved using brute-force methods (an Enumeration approach), DFS (an Exploration  
 181 approach), or Dijkstra’s algorithm (an Intuition approach).  
 182

183 To address this, we collect a variety of graph algorithm problems and propose an **LLM-as-judge** ap-  
 184 proach for determining the classification of each problem. Specifically, for each problem—comprising  
 185 a graph description and a task description expressed in natural language—we first generate 300 Erdős-  
 186 Rényi (ER) graphs (Erdős et al., 1960) to replace the original graph description, creating a diverse set  
 187 of problem instances. We then prompt LRM<sub>s</sub> with these instances, obtain their responses, and use  
 188 Qwen-2.5-72B (Team, 2024; Yang et al., 2024) as a judge to assign a reasoning category (detailed  
 189 prompts are provided in Appendix K.1). **To validate these labels, three postgraduate students with**  
 190 **competitive programming experience manually review a subset of instances and reach almost identical**  
 191 **judgments to Qwen-2.5-72B**. Finally, we select 9 graph tasks in which the algorithms used by LRM<sub>s</sub>  
 192 for reasoning are relatively unambiguous, thereby facilitating more straightforward classification. We  
 193 present in Appendix H.1 the distribution of algorithms used by representative LRM<sub>s</sub> across different  
 194 tasks, illustrating how each task aligns with our proposed reasoning taxonomy.  
 195

196 We now introduce the tasks within each reasoning taxonomy, along with the definition of each task.  
 197 Based on the optimal time complexity of each task under the corresponding algorithm of the category,  
 198 we classify the tasks into three levels: easy, medium, and hard. More specified task definitions and  
 the time complexity of the optimal algorithm for each task are given in Appendix F.1.

199 **Enumeration:** (1) Maximum Degree node (easy): Identify the node with the maximum  
 200 degree in a given undirected graph and output its degree. (2) Maximum Weight Triangle  
 201 (medium): Given an undirected graph where each node is assigned a weight, find a triangle (a set of  
 202 three nodes connected to each other) that has the maximum sum of node weights and output its weight  
 203 sum. (3) Maximum Clique Problem (MCP, hard): Given an undirected graph, identify a  
 204 subgraph in which all nodes are pairwise connected, and the subgraph contains the maximum number  
 205 of nodes. Then, output the size of the subgraph.  
 206

207 **Exploration:** (1) PathSum (easy): Given an undirected tree where each edge is assigned a  
 208 weight, output the number of paths from the root to the leaves for which the sum of the edge weights  
 209 exceeds a given threshold  $\tau$ . (2) Distance- $k$  (medium): Given an undirected graph and a  
 210 specific node  $v$ , find all other nodes that are within  $k$ -hops from  $v$  and output the number of qualified  
 211 nodes. (3) Diameter (hard): Identify the longest shortest path between any two nodes in a  
 212 given undirected graph and output the length of this path.  
 213

214 **Intuition:** (1) Maximum  $k$ -core (MKC, easy): Find a subgraph in an undirected graph where  
 215 each node has a degree of at least  $k$  and the subgraph contains the maximum number of nodes. Then,  
 output the size of the subgraph. (2) Minimum Spanning Tree (MST, medium): Given a  
 connected, undirected graph with edge weights, find a spanning tree that connects all nodes with the  
 minimum possible total edge weight and output the total edge weight. (3) Distance Threshold

216	Enumeration	Exploration	Intuition
217	 <p><b>Maximum Degree</b></p> <p>You are given an <b>undirected graph</b> representing the <b>Dublin street network</b>. You are required to <b>find the street in the network that has maximum connections to other streets</b>.</p> <ul style="list-style-type: none"> <li>- <b>Streets in the network:</b> Lissadel Avenue, Gallymore Road, Dolphin Road...</li> <li>- <b>Intersections in the network:</b> Lissadel Avenue and Benbulbin Road, Lissadel Avenue and Knocknarea Avenue...</li> </ul> <p>Please determine which street intersects with the greatest number of other streets, and <b>provide the total count of its intersections</b>.</p>	 <p><b>Distance-<math>k</math></b></p> <p>You are given an <b>undirected graph</b> representing the <b>London street network</b>. Given a street "Southerton Road", find all streets that are exactly distance 3 <b>away from the street</b>.</p> <ul style="list-style-type: none"> <li>- <b>Streets in the network:</b> Southerton Road, Beadon Road, Hammersmith Broadway, Kilmarsh Road...</li> <li>- <b>Intersections in the network:</b> Beadon Road and Cameron Road, Beadon Road and Glenthorne Road, Beadon Road and Hammersmith Broadway...</li> </ul> <p>Please find the streets in distance 3 from the street Southerton Road and <b>output the number of these streets</b>.</p>	 <p><b>Maximum <math>k</math>-core</b></p> <p>You are given an <b>undirected graph</b> representing the <b>Rome street network</b>. You are required to <b>find a maximum <math>k</math>-core (<math>k = 2</math>) in this network</b>.</p> <ul style="list-style-type: none"> <li>- <b>Streets in the network:</b> Via Udine, Viale Ippocrate, Via Catania...</li> <li>- <b>Intersections in the network:</b> Via Udine and Via Catania, Via Udine and Via Lucca...</li> </ul> <p>Please find the maximum 2-core in the network and <b>output the number of streets in this 2-core</b>.</p>

Figure 2: Illustrative problem description.

(hard): Given an undirected graph with positive edge weights, find the node that can reach the smallest number of other nodes via shortest paths of length at most a threshold, and output that node.

### 2.3 GRAPH COLLECTION

To measure the semantic understanding capabilities of LRM<sub>s</sub> and further avoid data contamination, we construct graphs and graph problems using nodes and edges with real-world semantics. We collect graph problems from the following five scenarios: (1) **DBLP (Ley, 2002)**: An academic collaboration network consisting of over 1.3 million nodes (authors) and 5 million edges (collaborations). This graph is used for the PathSum task. (2) **Street Network (Boeing, 2025)**: A collection of city transportation networks worldwide, where nodes represent streets and edges represent intersections between two streets. We use the transportation networks of London (containing 36,281 nodes and 83,662 edges, for Distance- $k$ ), Sydney (containing 30,291 nodes and 69,126 edges, for Minimum Spanning Tree), Dublin (containing 3,962 nodes and 7,000 edges, for Maximum Degree node), and Rome (containing 14,719 nodes and 28,087 edges, for Maximum  $k$ -core). (3) **OpenFlight (Open-Flights)**: An airport network consisting of 3,390 nodes (airports) and 19,166 edges (flight routes) weighted by geographical distance. This graph is assigned to the Distance Threshold task. (4)

**Wikipedia (Yin et al., 2017)**: A web graph of Wikipedia hyperlinks collected in September 2011, containing 1,791,489 nodes (representing articles) and 28,511,807 edges (representing hyperlinks between articles). This graph is used for the Maximum Weight Triangle and Maximum Clique Problem tasks. (5) **DBpedia (Bizer et al., 2009)**: A knowledge graph subset with over 1 million nodes (entities) and 7 million edges (relationships). This graph is assigned to Diameter task.

For each task, we employ random walk or BFS sampling methods to extract subgraphs of varying sizes from the corresponding real-world graphs. These subgraph sizes are divided into six categories: level-1 (8–15 nodes), level-2 (16–30 nodes), level-3 (31–50 nodes), level-4 (51–80 nodes), level-5 (81–120 nodes), and level-6 (121–160 nodes). In total, the dataset comprises 2,700 problem instances, with each size category containing 50 samples. Examples of problem descriptions for selected tasks are presented in Figure 2, while detailed statistics are provided in Table 2.

## 3 EVALUATION

In this section, we evaluate the capabilities of LRM<sub>s</sub> in solving graph algorithm problems and analyze the limitations in their reasoning. Our evaluation focuses on two research questions: **RQ1**: Do LRM<sub>s</sub> excel at long-context reasoning? **RQ2**: What leads to over-thinking of LRM<sub>s</sub> when solving graph algorithm problems? Details of the evaluation setup are provided in Appendix E.

### 3.1 EXPLORING LRM<sub>s</sub>’ LONG-CONTEXT REASONING ABILITY (RQ1)

With the rapid extension of context windows in modern LRM<sub>s</sub>, a central question is whether these models can truly perform *understanding and reasoning* over long contexts. Graph algorithm problems

270  
 271  
 272  
 273  
 274  
 275  
 Table 3: We evaluate models on six graph scales (Level-1 to Level-6). For each size, performance is  
 measured on three task taxonomies: Enumeration, Exploration, and Intuition. The evaluation metrics  
 are  $\text{cons@k}$  ( $\text{c@k}$ ) and  $\text{pass@k}$  ( $\text{p@k}$ ). We also report per-scale averages across tasks. Results for  
 Level-1 and Level-2 are presented in the Appendix H.2. For each model scale, the best results are  
 highlighted in **bold**, while the second-best results are underlined.

Models	Level-3 (31-50 nodes)								Level-4 (51-80 nodes)							
	Enumeration		Exploration		Intuition		Avg.		Enumeration		Exploration		Intuition		Avg.	
	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$
Skywork-OR1-7B-Preview	0.13	<b>0.58</b>	0.10	0.35	0.00	0.03	0.08	0.32	0.06	0.44	0.06	0.30	0.01	0.06	0.04	0.27
Light-R1-7B-DS	0.11	0.53	0.07	0.35	0.01	0.08	0.06	0.32	0.08	0.34	0.07	0.30	0.00	0.03	0.05	0.22
Distill-Qwen-7B	0.13	0.59	0.05	0.39	0.01	0.13	0.06	0.37	0.12	0.41	0.07	0.39	0.00	0.07	0.06	<b>0.29</b>
Qwen2.5-7B	0.11	0.39	0.10	0.21	0.00	0.03	0.07	0.21	0.09	0.36	0.11	0.29	0.03	0.07	<u>0.08</u>	0.24
OpenThinker-7B	0.19	0.61	0.07	0.43	0.01	0.13	0.09	<u>0.39</u>	0.03	0.41	0.04	0.32	0.01	0.03	0.03	0.25
Qwen3-8B-no-thinking	0.17	0.51	0.11	0.49	0.01	0.10	<u>0.10</u>	0.37	0.12	0.48	0.09	0.34	0.00	0.03	0.07	0.28
Qwen3-8B	0.79	0.99	0.69	0.93	0.40	0.69	<b>0.63</b>	<b>0.87</b>	0.44	0.83	0.28	0.63	0.09	0.26	<b>0.27</b>	<b>0.58</b>
GPT-OSS-20B	0.59	0.78	0.50	0.67	0.39	0.59	0.49	0.68	0.43	0.54	0.40	0.50	0.31	0.43	0.38	0.49
Light-R1-32B	0.71	0.97	0.64	0.94	0.35	0.59	0.57	0.83	0.43	0.87	0.38	0.76	0.06	0.31	0.29	0.65
Skywork-OR1-32B	0.87	0.99	0.77	0.99	0.51	0.75	<b>0.72</b>	<b>0.91</b>	0.54	0.87	0.51	0.84	0.12	0.42	0.39	0.71
Distill-Qwen-32B	0.63	0.89	0.46	0.81	0.21	0.51	0.43	0.74	0.42	0.80	0.23	0.60	0.03	0.31	0.23	0.57
Qwen2.5-32B	0.25	0.61	0.10	0.43	0.04	0.19	0.13	0.41	0.16	0.49	0.06	0.31	0.04	0.13	0.09	0.31
OpenThinker-32B	0.75	0.97	0.56	0.91	0.30	0.57	0.54	0.82	0.44	0.87	0.31	0.74	0.11	0.30	0.29	0.64
QWQ-32B	0.69	0.88	0.69	0.93	0.53	0.73	<u>0.69</u>	<u>0.88</u>	0.68	0.96	0.53	0.84	0.30	0.58	<b>0.50</b>	<b>0.79</b>
Qwen3-32B	0.79	0.99	0.69	0.93	0.40	0.69	0.63	0.87	0.52	0.89	0.51	0.87	0.26	0.59	0.43	0.78
Qwen3-32B-no-thinking	0.33	0.69	0.35	0.81	0.02	0.21	0.23	0.57	0.20	0.63	0.16	0.57	0.02	0.16	0.13	0.45
Llama-3.3-70B	0.10	0.36	0.13	0.39	0.01	0.04	0.08	0.26	0.04	0.27	0.00	0.21	0.00	0.01	0.01	0.16
GPT-OSS-120B	0.76	0.82	0.78	0.88	0.68	0.92	0.74	0.87	0.56	0.74	0.56	0.76	0.52	0.75	0.55	0.75
Qwen3-235B-A22B-Thinking	0.98	1.00	0.96	1.00	0.96	0.99	<b>0.97</b>	<b>1.00</b>	0.88	0.98	0.89	0.99	0.81	0.92	<b>0.86</b>	<b>0.96</b>
Qwen3-235B-A22B-Instruct	0.91	0.98	0.87	0.97	0.80	0.97	<u>0.86</u>	<u>0.97</u>	0.77	0.97	0.79	0.98	0.66	0.89	0.74	0.95
Models	Level-5 (81-120 nodes)								Level-6 (121-160 nodes)							
	Enumeration		Exploration		Intuition		Avg.		Enumeration		Exploration		Intuition		Avg.	
	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$	$\text{c@k}$	$\text{p@k}$
GPT-OSS-20B	0.51	0.68	0.18	0.33	0.16	0.27	0.28	0.43	0.28	0.48	0.10	0.14	0.04	0.09	0.14	0.24
Light-R1-32B	0.80	0.80	0.45	0.45	0.01	0.07	<b>0.42</b>	0.44	0.22	0.68	0.05	0.23	0.00	0.01	0.09	0.31
Skywork-OR1-32B	0.67	0.67	0.51	0.51	0.00	0.00	<u>0.39</u>	0.39	0.10	0.47	0.03	0.23	0.00	0.00	0.04	0.24
Distill-Qwen-32B	0.45	0.83	0.09	0.37	0.00	0.03	0.18	0.41	0.26	0.77	0.07	0.31	0.00	0.03	0.11	0.37
Qwen2.5-32B	0.15	0.53	0.05	0.22	0.02	0.11	0.07	0.29	0.21	0.48	0.03	0.19	0.01	0.01	0.08	0.23
OpenThinker-32B	0.37	0.76	0.14	0.45	0.02	0.21	0.18	0.47	0.27	0.68	0.07	0.33	0.00	0.02	0.11	0.34
QWQ-32B	0.61	0.95	0.25	0.63	0.09	0.25	0.32	<b>0.61</b>	0.47	0.87	0.13	0.40	0.09	0.21	<b>0.23</b>	<b>0.49</b>
Qwen3-32B	0.51	0.92	0.21	0.59	0.09	0.25	0.27	<u>0.59</u>	0.42	0.80	0.09	0.43	0.09	0.19	<u>0.20</u>	0.47
Qwen3-32B-no-thinking	0.15	0.61	0.07	0.28	0.01	0.07	0.08	0.32	0.13	0.49	0.07	0.23	0.01	0.07	0.07	0.27
Llama-3.3-70B	0.07	0.21	0.01	0.11	0.00	0.02	0.03	0.11	0.02	0.22	0.00	0.07	0.00	0.00	0.01	0.10
GPT-OSS-120B	0.51	0.67	0.41	0.51	0.22	0.28	0.38	0.49	0.37	0.54	0.19	0.33	0.11	0.16	0.22	0.34
Qwen3-235B-A22B-Thinking	0.86	0.96	0.52	0.82	0.37	0.49	<b>0.58</b>	<b>0.76</b>	0.80	0.93	0.28	0.52	0.26	0.36	<b>0.45</b>	<b>0.50</b>
Qwen3-235B-A22B-Instruct	0.76	0.92	0.41	0.68	0.26	0.38	<u>0.48</u>	<u>0.66</u>	0.51	0.79	0.28	0.59	0.21	0.26	<u>0.33</u>	<b>0.55</b>
GPT5-mini	0.66	0.72	0.53	0.63	0.02	0.03	0.40	0.46	0.77	0.86	0.36	0.47	0.10	0.18	<u>0.41</u>	<b>0.50</b>
Deepseek-V3	0.56	0.71	0.34	0.54	0.00	0.00	0.30	0.42	0.40	0.54	0.22	0.32	0.00	0.00	0.21	0.29
Deepseek-R1	0.82	0.82	0.69	0.69	0.15	0.15	<b>0.55</b>	<u>0.55</u>	0.79	0.79	0.50	0.50	0.05	0.05	<b>0.45</b>	0.45
Gemini-2.5-pro	0.73	0.90	0.46	0.63	0.17	0.30	<u>0.45</u>	<b>0.61</b>	0.63	0.86	0.37	0.54	0.08	0.1	0.36	<b>0.50</b>

312  
 313  
 314  
 315  
 316  
 317  
 318  
 319  
 320  
 321  
 322  
 323  
 provide an ideal testbed: they allow scalable control over input length, and their solutions cannot  
 be trivially extracted from the problem statement. To systematically **explore LRM**s’ **long-context**  
**reasoning ability**, we design several experiments. **First**, we vary the number of graph nodes,  
 and alternatively fix the graph structure while extending textual descriptions of nodes, to observe  
 performance changes as input length increases. **Second**, we collect and analyze erroneous responses  
 to identify the underlying causes of degraded performance in long-context reasoning.

324  
 325  
 326  
 327  
 328  
 329  
 330  
 331  
 332  
 333  
 334  
 335  
 336  
 337  
 338  
 339  
 340  
 341  
 342  
 343  
 344  
 345  
 346  
 347  
 348  
 349  
 350  
 351  
 352  
 353  
 354  
 355  
 356  
 357  
 358  
 359  
 360  
 361  
 362  
 363  
 364  
 365  
 366  
 367  
 368  
 369  
 370  
 371  
 372  
 373  
 374  
 375  
 376  
 377  
 378  
 379  
 380  
 381  
 382  
 383  
 384  
 385  
 386  
 387  
 388  
 389  
 390  
 391  
 392  
 393  
 394  
 395  
 396  
 397  
 398  
 399  
 400  
 401  
 402  
 403  
 404  
 405  
 406  
 407  
 408  
 409  
 410  
 411  
 412  
 413  
 414  
 415  
 416  
 417  
 418  
 419  
 420  
 421  
 422  
 423  
 424  
 425  
 426  
 427  
 428  
 429  
 430  
 431  
 432  
 433  
 434  
 435  
 436  
 437  
 438  
 439  
 440  
 441  
 442  
 443  
 444  
 445  
 446  
 447  
 448  
 449  
 450  
 451  
 452  
 453  
 454  
 455  
 456  
 457  
 458  
 459  
 460  
 461  
 462  
 463  
 464  
 465  
 466  
 467  
 468  
 469  
 470  
 471  
 472  
 473  
 474  
 475  
 476  
 477  
 478  
 479  
 480  
 481  
 482  
 483  
 484  
 485  
 486  
 487  
 488  
 489  
 490  
 491  
 492  
 493  
 494  
 495  
 496  
 497  
 498  
 499  
 500  
 501  
 502  
 503  
 504  
 505  
 506  
 507  
 508  
 509  
 510  
 511  
 512  
 513  
 514  
 515  
 516  
 517  
 518  
 519  
 520  
 521  
 522  
 523  
 524  
 525  
 526  
 527  
 528  
 529  
 530  
 531  
 532  
 533  
 534  
 535  
 536  
 537  
 538  
 539  
 540  
 541  
 542  
 543  
 544  
 545  
 546  
 547  
 548  
 549  
 550  
 551  
 552  
 553  
 554  
 555  
 556  
 557  
 558  
 559  
 560  
 561  
 562  
 563  
 564  
 565  
 566  
 567  
 568  
 569  
 570  
 571  
 572  
 573  
 574  
 575  
 576  
 577  
 578  
 579  
 580  
 581  
 582  
 583  
 584  
 585  
 586  
 587  
 588  
 589  
 590  
 591  
 592  
 593  
 594  
 595  
 596  
 597  
 598  
 599  
 600  
 601  
 602  
 603  
 604  
 605  
 606  
 607  
 608  
 609  
 610  
 611  
 612  
 613  
 614  
 615  
 616  
 617  
 618  
 619  
 620  
 621  
 622  
 623  
 624  
 625  
 626  
 627  
 628  
 629  
 630  
 631  
 632  
 633  
 634  
 635  
 636  
 637  
 638  
 639  
 640  
 641  
 642  
 643  
 644  
 645  
 646  
 647  
 648  
 649  
 650  
 651  
 652  
 653  
 654  
 655  
 656  
 657  
 658  
 659  
 660  
 661  
 662  
 663  
 664  
 665  
 666  
 667  
 668  
 669  
 670  
 671  
 672  
 673  
 674  
 675  
 676  
 677  
 678  
 679  
 680  
 681  
 682  
 683  
 684  
 685  
 686  
 687  
 688  
 689  
 690  
 691  
 692  
 693  
 694  
 695  
 696  
 697  
 698  
 699  
 700  
 701  
 702  
 703  
 704  
 705  
 706  
 707  
 708  
 709  
 710  
 711  
 712  
 713  
 714  
 715  
 716  
 717  
 718  
 719  
 720  
 721  
 722  
 723  
 724  
 725  
 726  
 727  
 728  
 729  
 7

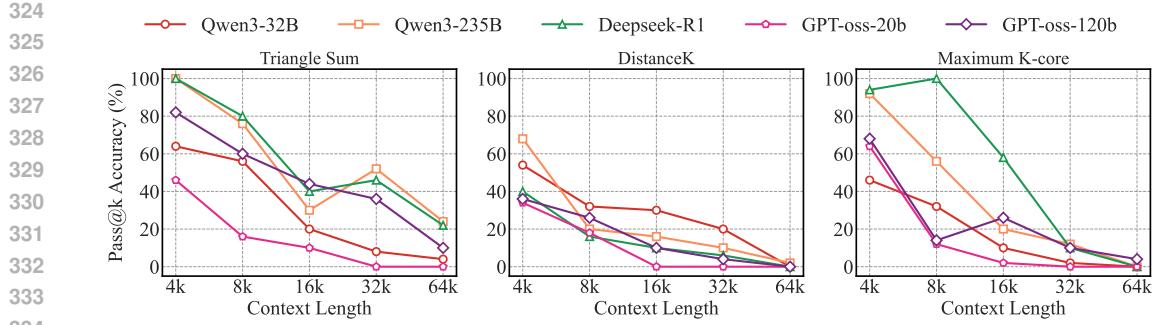


Figure 3: Models pass@k performance across different context length.

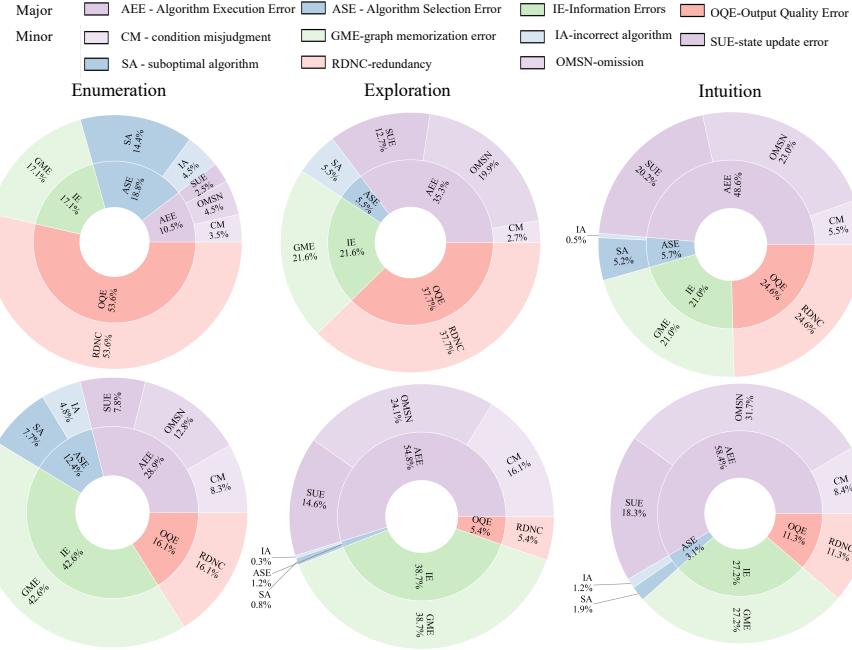


Figure 4: Error type distributions across reasoning taxonomies for Qwen3-32B (top) and its non-reasoning variant Qwen3-32B-no-thinking (bottom).

nodes and 200 edges each. By adjusting node name length, we generate five datasets with average token lengths of 4k, 8k, 16k, 32k, and 64k. We then evaluate five representative models—Qwen3-32B, Qwen3-235B-Thinking, GPT-oss-20B, GPT-oss-120B, and DeepSeek-R1—under these conditions, with results summarized in Figure 3.

**Error Analysis.** We examine the types of errors that LRM s commonly make when solving graph algorithm problems. Our methodology begins with establishing a taxonomy of **error categories** (see Appendix H.3), followed by the development of an automated **error detection pipeline**. Specifically, to support annotation and improve labeling accuracy, we reorganize LRM responses into coherent **paragraphs**. Specifically, we collect LRM s responses with incorrect final answers, segment them by “\n\n,” and index each paragraph as “ $\langle i \rangle$ ,” where  $i$  is the segment ID. We then use `qwen-2.5-72B` to merge paragraphs into **sections** by summarizing each response into tuples of (*start-index*, *end-index*, *summary*), which define the regrouped structure. Next, `O3-mini` annotates each section with **one or more error categories**, from which we compute the error-type distribution. In addition, three postgraduate students with competitive programming experience independently review a subset of annotations and confirm their consistency with the LLM-based labels. Detailed prompts and the full transformation process are provided in Appendix K.1 and K.2. Figure 4 presents the error distribution for Qwen3-32B and its non-reasoning counterpart across three taxonomies, with results for other models and detailed case studies in Appendices H.3 and K.3.

① **LRMs cannot effectively handle long-context reasoning.** Our two complementary experiments consistently reveal the limitations of LRM s in processing extended text inputs. As shown in Table 5, model accuracy decreases sharply as graph size grows. For example, Qwen3-32B achieves 87.0% pass@k accuracy on graphs with 31-50 nodes, but drops to 59.0% with 81-120 nodes and further to 47.0% beyond 121-160 nodes. Similar drops are observed across all evaluated models. While this decline could be attributed to increased structural complexity, additional experiments disentangle structural difficulty from textual length. By fixing the graph structure while gradually lengthening node descriptions, we observe a comparable pattern: on the TRIANGLE SUM task, DeepSeek-R1 achieves 100.0% at 4k tokens but only 22.0% at 64k tokens, while GPT-oss-120B falls from 68.0% to 10.0%. Taken together, these results point to a common bottleneck: as context length increases—whether from more graph nodes or longer textual descriptions—LRMs fail to maintain stable reasoning traces, resulting in significant performance degradation.

② **LRMs’ long-context reasoning is constrained by three key bottlenecks: coarse step-by-step execution, weak memory capability, and excessive redundancy.** First, LRM<sub>s</sub> often fail in algorithm execution even after selecting the correct strategy, as Algorithm Execution Errors (AEE) far exceed Algorithm Selection Errors (ASE)—for example, as shown in Figure 4, Qwen3-32B records AEE rates of 35.3% and 48.6% versus ASE rates of only 5.5% and 5.7% in *Exploration* and *Intuition*, respectively. The most frequent issues are State Update Errors and Omissions, indicating that models grasp high-level plans but falter on procedural details. Second, LRM<sub>s</sub> struggle to maintain an accurate representation of graph structure, leading to Graph Memorization Errors (GME): while non-reasoning variants reach error rates up to 42.6%, reasoning models such as Qwen3-32B still exhibit more than 21% GME in *Exploration* and *Intuition*. Finally, reasoning models frequently introduce excessive redundancy—for instance, Qwen3-32B shows 37.7% redundancy errors in *Exploration*—where verbose and repetitive steps inflate reasoning traces without improving accuracy, and sometimes even cause additional confusion. Importantly, such redundancy reflects a form of *over-thinking*, which we will further analyze in the next section.

③ **Both reasoning and non-reasoning models achieve the best performance on *Enumeration* tasks, followed by *Exploration*, and the worst on *Intuition*.** This observation can be explained from two perspectives. First, within each node-size level, almost all models follow the pattern *Enumeration* > *Exploration* > *Intuition*. For instance, as shown at Table 3, at Level-3, GPT-OSS-20B obtains *cons@k* scores of 0.59, 0.50, and 0.39 on *Enumeration*, *Exploration*, and *Intuition*, respectively. Second, considering how large a graph each model can handle, we find that some remain more robust on *Enumeration*. For example, at Level-6, GPT-5-mini reaches 0.77 *cons@k* accuracy on *Enumeration*, but drops below 0.4 on *Exploration* and to only about 0.1 on *Intuition*. This indicates that LRMs are relatively competent at systematic, enumeration-based reasoning but struggle substantially with intuition-driven reasoning. To further support this finding, we provide a normalized analysis using the Z-score metric in Appendix H.2.

### 3.2 EXPLORING UNDERLYING MECHANISMS OF LRM<sup>S</sup> OVER-THINKING (RQ2)

In error analysis, we observe that reasoning models often produce substantial redundancy when solving graph problems, with responses containing unnecessary information. This phenomenon is closely related to the *over-thinking* behavior characteristic of LRM s, where the model inefficiently expends resources by engaging in excessive verification or redundant checks even after reaching a final answer (Lu et al., 2025; Chen et al., 2024c). Such over-thinking not only generates superfluous tokens, but also leads to wasted computational resources and longer response times. In this section, we investigate LRM s' over-thinking phenomenon in graph algorithm problem solving from two perspectives: *post-answer generation* and *self-verification*.

**429 Response Partitioning via High-Entropy Tokens.** To facilitate subsequent evaluation, we parti-  
430 tion LRM-generated responses into segments. Our partitioning approach is guided by token-level  
431 generation entropy, a metric that quantifies the model’s uncertainty at each step of the generation  
process (See Appendix F.3 for more details). High entropy values often coincide with moments

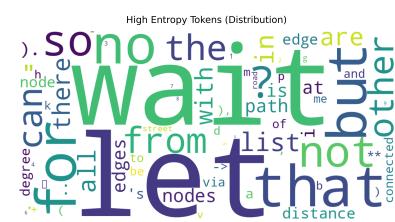


Figure 5: High-entropy tokens in LRM inference traces

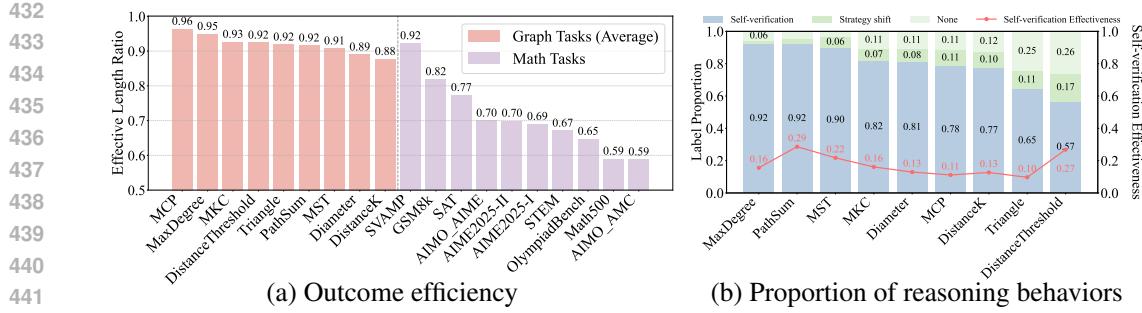


Figure 6: Overthinking analysis of QWQ-32B (See Appendix H.4 for more results).

of self-verification or strategy-shifting, which may indicate over-thinking (Wang et al., 2025). By computing this token-level entropy across reasoning traces and visualizing it with a word cloud (Figure 5), we find that tokens such as “*wait*”, “*but*”, and “*so*” consistently exhibit high entropy. This observation aligns with prior studies (Li et al., 2025a; Lu et al., 2025), which have identified these tokens as typical triggers of over-thinking. Therefore, we adopt these tokens as **special partition markers**. Formally, let a response be a token sequence  $T = (w_1, w_2, \dots, w_m)$ . If special tokens appear at positions  $i_1, i_2, \dots, i_n$ , then  $T$  is partitioned into  $n + 1$  consecutive **segments**:

$$T \Rightarrow (w_1, \dots, w_{i_1-1}), (w_{i_1+1}, \dots, w_{i_2-1}), \dots, (w_{i_n+1}, \dots, w_m). \quad (1)$$

**Post-Answer Generation.** One possible source of over-thinking is **post-answer generation**, i.e., the model keeps producing tokens even after outputting a correct answer (Sui et al., 2025; Chen et al., 2024c; 2025). To quantify this behavior, we adopt the **outcome efficiency** metric (Chen et al., 2024c), the fraction of tokens generated up to the first correct answer relative to the full response length:

$$\zeta_O = \frac{1}{N} \sum_{i=1}^N \frac{\hat{T}_i}{T_i}, \quad (2)$$

where  $N$  is the number of responses containing the correct answer,  $T_i$  the total response length, and  $\hat{T}_i$  the tokens up to the first correct answer. We compute  $\hat{T}_i$  by prompting Qwen2.5-72B with the original problem, the segmented LRM response, and the ground-truth answer, asking it to locate the first correct segment (Prompts can be seen in Appendix K.1). Outcome efficiency of QWQ-32B on graph and math problems is shown in Figure 6 (a).

**Self-Verification.** Another relevant behavior is **self-verification**, where LRMs attempt to re-check earlier reasoning or conclusions. To analyze this process, we again use Qwen2.5-72B: for each segment  $i$ , the model classifies it (given all prior segments) as *self-verification*, *strategy-shift*, or *other*. If labeled as self-verification, the segment is further judged as *effective* when it correctly identifies prior errors. Detailed prompts are listed in Appendix K.1, and Figure 6(b) summarizes the distribution of segment types and their effectiveness.

**④ Frequent yet ineffective self-verification is a primary driver of over-thinking in LRMs when solving graph algorithm problems.** As shown in Figure 6, outcome efficiency (a) remains consistently high on graph tasks—above 0.88 in all cases—and is higher than on math tasks. This indicates that LRMs seldom continue generating tokens long after reaching the correct answer, suggesting that post-answer generation is not the main source of over-thinking in graph reasoning. In contrast, panel (b) shows that self-verification occupies a substantial fraction of reasoning traces across different tasks, yet its effectiveness is strikingly low: in most cases fewer than 30% of self-verification attempts successfully detect prior errors. Instead, the majority of these segments simply restate earlier steps or add irrelevant content, inflating response length without improving correctness. These observations together point to frequent but ineffective self-verification as the main mechanism underlying the over-thinking phenomenon in LRMs when solving graph algorithm problems.

### 3.3 DISCUSSION

We further discuss two aspects: (1) the potential benefits of having LRMs generate code and invoke external tools to solve graph algorithm problems, and (2) the extent to which the insights derived from our benchmark generalize beyond the specific tasks and settings considered in this work.

Table 4: Model performance (pass@k / cons@k (%)) with and without code-based tool usage.

Task	Metric	Level-5 (81–120 nodes)			Level-6 (121–160 nodes)		
		Owen3-32B	Gemini-2.5-pro	GPT5-mini	Owen3-32B	Gemini-2.5-pro	GPT5-mini
Triangle	Base	76.0 / 66.6	90.0 / 80.0	56.7 / 50.0	60.0 / 42.2	86.7 / 83.3	86.7 / 70.0
	Code	<b>96.0 / 93.3</b>	<b>94.0 / 92.0</b>	<b>99.0 / 99.0</b>	<b>92.0 / 88.0</b>	<b>90.0 / 88.0</b>	<b>99.0 / 99.0</b>
DistanceK	Base	40.0 / 20.0	43.3 / 30.0	60.0 / 53.3	36.0 / 10.0	43.3 / 33.3	26.7 / 13.3
	Code	<b>98.0 / 96.0</b>	<b>100.0 / 100.0</b>	<b>100.0 / 100.0</b>	<b>96.7 / 96.7</b>	<b>100.0 / 100.0</b>	<b>100.0 / 100.0</b>
MST	Base	0 / 0	16.6 / 3.3	3.3 / 0	0 / 0	0 / 0	0 / 0
	Code	<b>97.3 / 94.7</b>	<b>98.7 / 98.0</b>	<b>100.0 / 100.0</b>	<b>88.0 / 84.0</b>	<b>92.0 / 90.0</b>	<b>100.0 / 100.0</b>

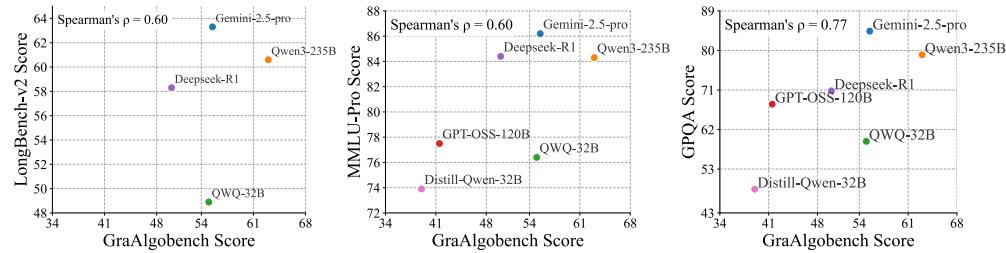


Figure 7: Model-wise correlation with external benchmarks .

**Solving Graph Algorithms Problems by External Tools.** To better understand the impact of tool usage on our benchmark, we conduct an additional experiment where models are explicitly instructed to solve problems by writing code and then using a code-execution tool (examples are given in appendix G). Concretely, we compare three LRMs of different strengths (Owen3-32B, Gemini-2.5-pro, GPT5-mini) on three representative tasks (Triangle, DistanceK, MST), at graph size level-5 and level-6. We report pass@k and cons@k with and without tool calls (See Table 4).

Across models and tasks, enabling code and tool usage consistently and substantially improves performance, especially on the most challenging instances. However, delegating the hardest parts of the problem—such as graph memorization and step-by-step algorithmic execution—to external tools no longer directly evaluates the model’s intrinsic reasoning ability. For this reason, we do not adopt the code-generation setting as our primary evaluation protocol.

**Conclusion Generalization.** To verify that our findings generalize beyond graph-specific scenarios, we conduct two additional experiments: (1) we measure model-wise correlations between GraAlgobench and established benchmarks spanning long-context and general reasoning—LongBench-v2 (Bai et al., 2025), GPQA (Rein et al., 2024), and MMLU-Pro (Wang et al., 2024) (Figure 7); and (2) we vary the graph serialization format (edge lists, adjacency lists, Markdown tables, and numeric node IDs; examples in Appendix G) and compare performance across formats (Table 6 in Appendix I).

According to standard guidelines for interpreting correlation coefficients (Mukaka, 2012), values in the range 0.50–0.79 indicate a moderate to strong association. The observed correlations between GraAlgobench and GPQA ( $\rho = 0.77$ ) and between GraAlgobench and LongBench-v2 ( $\rho = 0.60$ ) therefore suggest that the *complex reasoning skills* assessed by our benchmark align closely with those underlying high-level general reasoning and long-context understanding. In addition, across graph sizes, models achieve *similar performance under all four representation formats*, indicating that performance is *largely independent of input format* and primarily reflects intrinsic reasoning ability rather than superficial properties of the graph encoding.

## 4 CONCLUSION

We introduced GRALGOBENCH, a benchmark that leverages graph algorithm problems to probe reasoning in LRMs with effective long-context evaluation, fine-grained difficulty control, and programmatic assessment. Our results highlight two central limitations: accuracy drops substantially as context length grows, and over-thinking emerges from frequent yet ineffective self-verification. By uncovering these challenges, GRALGOBENCH establishes graph algorithms as a practical and rigorous foundation for advancing the robustness and efficiency of future reasoning models.

## 540 REFERENCES

541

542 Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K  
543 Arora, Yu Bai, Bowen Baker, Haiming Bao, et al. gpt-oss-120b & gpt-oss-20b model card. *arXiv*  
544 *preprint arXiv:2508.10925*, 2025.

545 Sara Ahmadian, Sreenivas Gollapudi, Gregory Hutchins, Kostas Kollias, and Xizhi Tan. Extracting  
546 small subgraphs in road networks. In *Proceedings of the ACM Web Conference 2024*, pp. 493–502,  
547 2024.

548 Yushi Bai, Shangqing Tu, Jiajie Zhang, Hao Peng, Xiaozhi Wang, Xin Lv, Shulin Cao, Jiazheng Xu,  
549 Lei Hou, Yuxiao Dong, et al. Longbench v2: Towards deeper understanding and reasoning on  
550 realistic long-context multitasks. In *Proceedings of the 63rd Annual Meeting of the Association for*  
551 *Computational Linguistics (Volume 1: Long Papers)*, pp. 3639–3664, 2025.

552

553 Christian Bizer, Jens Lehmann, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak,  
554 and Sebastian Hellmann. Dbpedia-a crystallization point for the web of data. *Journal of web*  
555 *semantics*, 7(3):154–165, 2009.

556 Geoff Boeing. Modeling and analyzing urban networks and amenities with osmnx. *Geographical*  
557 *Analysis*, 2025.

558

559 Matteo Ceccarello, Andrea Pietracaprina, and Geppino Pucci. Fast and accurate fair k-center  
560 clustering in doubling metrics. In *Proceedings of the ACM Web Conference 2024*, pp. 756–767,  
561 2024.

562 Nuo Chen, Yuhan Li, Jianheng Tang, and Jia Li. Graphwiz: An instruction-following language model  
563 for graph computational problems. In *Proceedings of the 30th ACM SIGKDD Conference on*  
564 *Knowledge Discovery and Data Mining*, pp. 353–364, 2024a.

565

566 Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu,  
567 Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-  
568 thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*, 2025.

569

570 Xiaolong Chen, Yifan Song, and Jing Tang. Link recommendation to augment influence diffusion  
571 with provable guarantees. In *Proceedings of the ACM Web Conference 2024*, pp. 2509–2518,  
572 2024b.

573

574 Xingyu Chen, Jiahao Xu, Tian Liang, Zhiwei He, Jianhui Pang, Dian Yu, Linfeng Song, Qiuzhi Liu,  
575 Mengfei Zhou, Zhuosheng Zhang, et al. Do not think that much for  $2+3=?$  on the overthinking of  
576 o1-like llms. *arXiv preprint arXiv:2412.21187*, 2024c.

577

578 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit  
579 Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier  
with advanced reasoning, multimodality, long context, and next generation agentic capabilities.  
*arXiv preprint arXiv:2507.06261*, 2025.

580

581 Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to*  
*algorithms*. MIT press, 2022.

582

583 Jianbo Dai, Jianqiao Lu, Yunlong Feng, Guangtao Zeng, Rongju Ruan, Ming Cheng, Dong Huang,  
584 Haochen Tan, and Zhijiang Guo. Mhpp: Exploring the capabilities and limitations of language  
585 models beyond basic code generation. *arXiv preprint arXiv:2405.11430*, 2024.

586

587 Guanting Dong, Yifei Chen, Xiaoxi Li, Jiajie Jin, Hongjin Qian, Yutao Zhu, Hangyu Mao, Guorui  
588 Zhou, Zhicheng Dou, and Ji-Rong Wen. Tool-star: Empowering llm-brained multi-tool reasoner  
589 via reinforcement learning. *arXiv preprint arXiv:2505.16410*, 2025.

590

591 Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha  
592 Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models.  
*arXiv e-prints*, pp. arXiv–2407, 2024.

593

594 Paul Erdős, Alfréd Rényi, et al. On the evolution of random graphs. *Publ. math. inst. hung. acad. sci.*,  
5(1):17–60, 1960.

594 Bahare Fatemi, Jonathan Halcrow, and Bryan Perozzi. Talk like a graph: Encoding graphs for large  
 595 language models. *arXiv preprint arXiv:2310.04560*, 2023.

596

597 Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang,  
 598 Jinxin Chi, and Wanjun Zhong. Retool: Reinforcement learning for strategic tool use in llms.  
 599 *arXiv preprint arXiv:2504.11536*, 2025.

600 Elliot Glazer, Ege Erdil, Tamay Besiroglu, Diego Chicharro, Evan Chen, Alex Gunning, Caroline Falk-  
 601 man Olsson, Jean-Stanislas Denain, Anson Ho, Emily de Oliveira Santos, et al. Frontiernmath: A  
 602 benchmark for evaluating advanced mathematical reasoning in ai. *arXiv preprint arXiv:2411.04872*,  
 603 2024.

604

605 Aryan Gulati, Brando Miranda, Eric Chen, Emily Xia, Kai Fronsdal, Bruno de Moraes Dumont,  
 606 and Sanmi Koyejo. Putnam-axiom: A functional & static benchmark for measuring higher level  
 607 mathematical reasoning in llms. In *Forty-second International Conference on Machine Learning*,  
 608 2024.

609

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

613

614 Jiayan Guo, Lun Du, Hengyu Liu, Mengyu Zhou, Xinyi He, and Shi Han. Gpt4graph: Can large  
 615 language models understand graph structured data? an empirical evaluation and benchmarking.  
 616 *arXiv preprint arXiv:2305.15066*, 2023.

617

618 Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu,  
 619 Xu Han, Yujie Huang, Yuxiang Zhang, et al. Olympiadbench: A challenging benchmark for  
 620 promoting agi with olympiad-level bilingual multimodal scientific problems. *arXiv preprint  
 621 arXiv:2402.14008*, 2024.

622

623 Jujie He, Jiacai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang  
 624 Zhang, Jiacheng Xu, Wei Shen, et al. Skywork open reasoner 1 technical report. *arXiv preprint  
 625 arXiv:2505.22312*, 2025a.

626

627 Junda He, Jieke Shi, Terry Yue Zhuo, Christoph Treude, Jiamou Sun, Zhenchang Xing, Xiaoning  
 628 Du, and David Lo. From code to courtroom: Llms as the new software judges. *arXiv preprint  
 629 arXiv:2503.02246*, 2025b.

630

631 Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle  
 632 Cai, Hui Yuan, Runzhe Wang, et al. Math-perturb: Benchmarking llms' math reasoning abilities  
 633 against hard perturbations. *arXiv preprint arXiv:2502.06453*, 2025.

634

635 Zhen Huang, Zengzhi Wang, Shijie Xia, Xuefeng Li, Haoyang Zou, Ruijie Xu, Run-Ze Fan, Lyuman-  
 636 shan Ye, Ethan Chern, Yixin Ye, et al. Olympicarena: Benchmarking multi-discipline cognitive  
 637 reasoning for superintelligent ai. *Advances in Neural Information Processing Systems*, 37:19209–  
 638 19253, 2024.

639

640 Michael Ley. The dblp computer science bibliography: Evolution, research issues, perspectives. In  
 641 *International symposium on string processing and information retrieval*, pp. 1–10. Springer, 2002.

642

643 Chengpeng Li, Mingfeng Xue, Zhenru Zhang, Jiaxi Yang, Beichen Zhang, Xiang Wang, Bowen  
 644 Yu, Binyuan Hui, Junyang Lin, and Dayiheng Liu. Start: Self-taught reasoner with tools. *arXiv  
 645 preprint arXiv:2503.04625*, 2025a.

646

647 Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittweiser, Rémi Leblond, Tom  
 648 Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation  
 649 with alphacode. *Science*, 378(6624):1092–1097, 2022.

650

651 Yunxin Li, Baotian Hu, Haoyuan Shi, Wei Wang, Longyue Wang, and Min Zhang. Visiongraph:  
 652 Leveraging large multimodal models for graph theory problems in visual context. *arXiv preprint  
 653 arXiv:2405.04950*, 2024.

654

655 Zhiyuan Li, Yi Chang, and Yuan Wu. Think-bench: Evaluating thinking efficiency and chain-of-  
 656 thought quality of large reasoning models. *arXiv preprint arXiv:2505.22113*, 2025b.

648 Bill Yuchen Lin, Ronan Le Bras, Kyle Richardson, Ashish Sabharwal, Radha Poovendran, Peter  
 649 Clark, and Yejin Choi. Zebralogic: On the scaling limits of llms for logical reasoning. *arXiv*  
 650 *preprint arXiv:2502.01100*, 2025.

651 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,  
 652 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*  
 653 *arXiv:2412.19437*, 2024.

654 Ximing Lu, Seungju Han, David Acuna, Hyunwoo Kim, Jaehun Jung, Shrimai Prabhumoye, Niklas  
 655 Muennighoff, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro, et al. Retro-search:  
 656 Exploring untaken paths for deeper and efficient reasoning. *arXiv preprint arXiv:2504.04383*,  
 657 2025.

658 Zihan Luo, Xiran Song, Hong Huang, Jianxun Lian, Chenhao Zhang, Jinqi Jiang, and Xing Xie.  
 659 Graphinstruct: Empowering large language models with graph understanding and reasoning  
 660 capability. *arXiv preprint arXiv:2403.04483*, 2024.

661 Yingqian Min, Zhipeng Chen, Jinhao Jiang, Jie Chen, Jia Deng, Yiwen Hu, Yiru Tang, Jiapeng Wang,  
 662 Xiaoxue Cheng, Huatong Song, et al. Imitate, explore, and self-improve: A reproduction report on  
 663 slow-thinking reasoning systems. *arXiv preprint arXiv:2412.09413*, 2024.

664 Atsushi Miyauchi, Lorenzo Severini, and Francesco Bonchi. Local centrality minimization with  
 665 quality guarantees. In *Proceedings of the ACM Web Conference 2024*, pp. 410–421, 2024.

666 Mavuto M Mukaka. A guide to appropriate use of correlation coefficient in medical research. *Malawi*  
 667 *medical journal*, 24(3):69–71, 2012.

668 Lutz Oettershagen, Honglian Wang, and Aristides Gionis. Finding densest subgraphs with edge-color  
 669 constraints. In *Proceedings of the ACM Web Conference 2024*, pp. 936–947, 2024.

670 OpenAI. Graphwalks. <https://huggingface.co/datasets/openai/graphwalks>,  
 671 2025.

672 OpenAI. Introducing gpt-5. *OpenAI*, August 2025. Accessed: 2025-08-07.

673 OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>, 2025. Accessed: 15 March 2025. Referenced on  
 674 pp. 1, 4, 6.

675 OpenFlights. <https://openflights.org/>. Accessed: 2024-05-25.

676 Jiayang Pang, Chenhao Ma, and Yixiang Fang. A similarity-based approach for efficient large  
 677 quasi-clique detection. In *Proceedings of the ACM Web Conference 2024*, pp. 401–409, 2024.

678 Ivo Petrov, Jasper Dekoninck, Lyuben Baltadzhiev, Maria Drencheva, Kristian Minchev, Mislav  
 679 Balunović, Nikola Jovanović, and Martin Vechev. Proof or bluff? evaluating llms on 2025 usa  
 680 math olympiad. *arXiv preprint arXiv:2503.21934*, 2025.

681 Yiwei Qin, Xuefeng Li, Haoyang Zou, Yixiu Liu, Shijie Xia, Zhen Huang, Yixin Ye, Weizhe Yuan,  
 682 Hector Liu, Yuanzhi Li, et al. O1 replication journey: A strategic progress report–part 1. *arXiv*  
 683 *preprint arXiv:2410.18982*, 2024.

684 David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani,  
 685 Julian Michael, and Samuel R Bowman. Gpqa: A graduate-level google-proof q&a benchmark. In  
 686 *First Conference on Language Modeling*, 2024.

687 Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu,  
 688 Andrew Wen, Shaochen Zhong, Hanjie Chen, et al. Stop overthinking: A survey on efficient  
 689 reasoning for large language models. *arXiv preprint arXiv:2503.16419*, 2025.

690 Haoxiang Sun, Yingqian Min, Zhipeng Chen, Wayne Xin Zhao, Lei Fang, Zheng Liu, Zhongyuan  
 691 Wang, and Ji-Rong Wen. Challenging the boundaries of reasoning: An olympiad-level math  
 692 benchmark for large language models. *arXiv preprint arXiv:2503.21380*, 2025.

702 Mirac Suzgun, Nathan Scales, Nathanael Schärlí, Sebastian Gehrmann, Yi Tay, Hyung Won Chung,  
 703 Akanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks  
 704 and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.

705

706 Jianheng Tang, Qifan Zhang, Yuhan Li, and Jia Li. Grapharena: Benchmarking large language  
 707 models on graph computational problems. *arXiv e-prints*, pp. arXiv–2407, 2024.

708

709 OpenThoughts Team. Open Thoughts. <https://open-thoughts.ai>, January 2025a.

710

711 Qwen Team. Qwen2.5: A party of foundation models, September 2024. URL <https://qwenlm.github.io/blog/qwen2.5/>.

712

713 Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, 2025b.

714

715 Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov.  
 716 Can language models solve graph problems in natural language? *Advances in Neural Information  
 717 Processing Systems*, 36:30840–30861, 2023.

718

719 Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen,  
 720 Jianxin Yang, Zhenru Zhang, et al. Beyond the 80/20 rule: High-entropy minority tokens drive  
 721 effective reinforcement learning for llm reasoning. *arXiv preprint arXiv:2506.01939*, 2025.

722

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

726

727 Anjiang Wei, Jiannan Cao, Ran Li, Hongyu Chen, Yuhui Zhang, Ziheng Wang, Yuan Liu, Thi-  
 728 ago SFX Teixeira, Diyi Yang, Ke Wang, et al. Equibench: Benchmarking large language models’  
 729 understanding of program semantics via equivalence checking. *arXiv preprint arXiv:2502.12466*,  
 2025.

730

731 Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu  
 732 Tang, Xiaowei Lv, et al. Light-r1: Curriculum sft, dpo and rl for long cot from scratch and beyond.  
 733 *arXiv preprint arXiv:2503.10460*, 2025.

734

735 Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-  
 736 Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, et al. Livebench: A challenging, contamination-  
 737 limited llm benchmark. *arXiv preprint arXiv:2406.19314*, 2024.

738

739 Hao Xu, Xiangru Jian, Xinjian Zhao, Wei Pang, Chao Zhang, Suyuchen Wang, Qixin Zhang,  
 740 Zhengyuan Dong, Joao Monteiro, Bang Liu, et al. Graphomni: A comprehensive and extend-  
 741 able benchmark framework for large language models on graph-theoretic tasks. *arXiv preprint  
 742 arXiv:2504.12764*, 2025.

743

744 An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li,  
 745 Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang,  
 746 Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai,  
 747 Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng  
 748 Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai  
 749 Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan  
 750 Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang  
 751 Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. Qwen2  
 752 technical report. *arXiv preprint arXiv:2407.10671*, 2024.

753

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

756

757 Lei Yang, Renren Jin, Ling Shi, Jianxiang Peng, Yue Chen, and Deyi Xiong. Probench: Benchmarking  
 758 large language models in competitive programming. *arXiv preprint arXiv:2502.20868*, 2025b.

756 Yue Yang, MingKang Chen, Qihua Liu, Mengkang Hu, Qiguang Chen, Gengrui Zhang, Shuyue  
 757 Hu, Guangtao Zhai, Yu Qiao, Yu Wang, et al. Truly assessing fluid intelligence of large language  
 758 models through dynamic reasoning evaluation. *arXiv preprint arXiv:2506.02648*, 2025c.

759

760 Hao Yin, Austin R Benson, Jure Leskovec, and David F Gleich. Local higher-order graph clustering.  
 761 In *Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and*  
 762 *data mining*, pp. 555–564, 2017.

763

764 Jifan Yu, Xiaozhi Wang, Shangqing Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun  
 765 Yao, Xiaohan Zhang, Hanming Li, et al. Kola: Carefully benchmarking world knowledge of large  
 766 language models. *arXiv preprint arXiv:2306.09296*, 2023.

767

768 Zhaojian Yu, Yilun Zhao, Arman Cohan, and Xiao-Ping Zhang. Humaneval pro and mbpp pro: Evalu-  
 769 uating large language models on self-invoking code generation. *arXiv preprint arXiv:2412.21199*,  
 770 2024.

771

772 Zike Yuan, Ming Liu, Hui Wang, and Bing Qin. Gracore: Benchmarking graph comprehension and  
 773 complex reasoning in large language models. *arXiv preprint arXiv:2407.02936*, 2024.

774

775 Qifan Zhang, Nuo Chen, Zehua Li, Miao Peng, Jing Tang, and Jia Li. Improving llms’ generalized  
 776 reasoning abilities by graph problems. *arXiv preprint arXiv:2507.17168*, 2025.

777

778 Zeyang Zhang, Xin Wang, Ziwei Zhang, Haoyang Li, Yijian Qin, and Wenwu Zhu. Llm4dyg: can  
 779 large language models solve spatial-temporal problems on dynamic graphs? In *Proceedings of the*  
 780 *30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4350–4361, 2024.

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810 A ETHIC STATEMENT  
811812 This work focuses on benchmarking large reasoning models through graph algorithm problems. The  
813 study does not involve human subjects, personally identifiable information, or sensitive data. All  
814 experiments are conducted using publicly available models and synthetic datasets. We have taken  
815 care to ensure that the research does not promote harmful applications and is intended solely for  
816 advancing understanding of reasoning in AI systems.  
817818 B REPRODUCTIVITY STATEMENT  
819820 We provide detailed task definitions, dataset generation procedures, and evaluation protocols to ensure  
821 full reproducibility. All experiments are conducted with publicly available models and standard-  
822 ized settings, as described in Section E. The complete codebase, including dataset generators and  
823 evaluation scripts, is released at the project repository to facilitate replication and future extensions.  
824825 C THE USE OF LARGE LANGUAGE MODELS (LLMs)  
826827 During the preparation of this work, we made use of LLMs in several stages of writing and experi-  
828 mentation:829

- **Language Refinement:** LLMs (such as OpenAI’s ChatGPT) were employed to improve the  
830 readability and academic style of the manuscript, helping to enhance clarity and consistency.  
831
- **Evaluation Assistance:** We employed LLMs as judges to help with evaluation. These  
832 models supported the assessment process and complemented our programmatic evaluation  
833 pipeline.  
834
- **Programming Support:** The Cursor environment, which incorporates AI assistance, was  
835 used for generating, modifying, and refactoring segments of the codebase. This contributed  
836 to faster implementation and better code quality.  
837

838 D RELATED WORKS  
839840 **Benchmarking LLMs on Graph Algorithm Problems.** Graph algorithm problems demand a  
841 deep understanding of structural information and long-range multi-step reasoning, making them  
842 particularly challenging for LLMs. For this reason, they have been widely used in prior studies  
843 to evaluate LLM capabilities. NLGraph (Wang et al., 2023), GPT4Graph (Guo et al., 2023), and  
844 GraphQA (Fatemi et al., 2023) are among the earliest works, assessing LLMs on graph algorithm  
845 problems by encoding tasks with LLMs alone or with a GNN–LLM combination. Building on  
846 this line of research, LLM4DyG (Zhang et al., 2024) examines tasks on dynamic graphs, while  
847 GraphInstruct (Luo et al., 2024) provides a broader benchmark with 21 tasks ranging from node-level  
848 to graph-level. GraphArena (Tang et al., 2024) evaluates LLMs on real-world graphs with more  
849 fine-grained metrics, and VisionGraph (Li et al., 2024) explores their capabilities on image graphs.  
850 Most recently, GraCore (Yuan et al., 2024) evaluates LLMs from the perspectives of graph reasoning  
851 and graph understanding, and GraphOmni (Xu et al., 2025) investigates their performance across  
852 diverse graph types, serialization formats, and prompting schemes. However, none of these works  
853 evaluate the performance of O1-like large reasoning models (LRMs) or analyze the challenges LRM  
854 face when solving graph algorithm problems, which motivates our study.  
855856 **Evaluating Large Reasoning Models.** The emergence of O1-like models, equipped with long  
857 chains of thought and advanced reasoning strategies such as self-verification, strategy-shift, and  
858 backtracking, has substantially advanced the reasoning capabilities of LLMs. This progress has  
859 motivated extensive efforts to evaluate LRM across diverse domains (Petrov et al., 2025; He et al.,  
860 2024; Huang et al., 2024; Gulati et al., 2024; Glazer et al., 2024; Li et al., 2025b; Huang et al., 2025;  
861 Chen et al., 2024c; Wang et al., 2023; Li et al., 2022; Dai et al., 2024; Yang et al., 2025b; Yu et al.,  
862 2024; Wei et al., 2025; He et al., 2025b; Zheng et al., 2025; White et al., 2024; Suzgun et al., 2022; Lin  
863 et al., 2025). Among these, mathematics has become the predominant benchmarking arena, yielding  
864 many new insights into the strengths and weaknesses of LRM. For instance, Putnam-AXIOM (Gulati

864 et al., 2024), FrontierMath (Glazer et al., 2024), and OlympiadBench (Sun et al., 2025) introduce  
 865 increasingly challenging problems to probe the reasoning limits of models, while other works (Chen  
 866 et al., 2024c; Wang et al., 2023) analyze phenomena such as *over-thinking* and *under-thinking* and  
 867 explore potential mitigations. However, no prior work has systematically evaluated LRM<sub>s</sub> on *graph*  
 868 *algorithm problems*, which demand structural reasoning and memory—skills missing from existing  
 869 benchmarks. To close this gap, we present **GRALGOBENCH**, the first graph-based benchmark for  
 870 rigorous and application-relevant LRM<sub>s</sub> evaluation.

## 872 E EXPERIMENTAL SETUP

874 **Models:** We evaluate a broad range of reasoning models as well as non-reasoning models in order to  
 875 ensure comprehensive coverage. For open-source models, we include the Qwen3 series (Yang et al.,  
 876 2025a), QwQ (Team, 2025b), Qwen-2.5 series (Team, 2024), OpenThinker2 series (Team, 2025a),  
 877 Skywork series (He et al., 2025a), Light-R1 series (Wen et al., 2025), Llama-3.3 series (Dubey et al.,  
 878 2024), GPT-OSS series (Agarwal et al., 2025), Deepseek-V3 (Liu et al., 2024), and Deepseek-R1  
 879 series (Guo et al., 2025). For closed-source models, we evaluate GPT5-mini (OpenAI, 2025) and  
 880 Gemini-2.5-pro (Comanici et al., 2025).

881 **Evaluation Metrics.** Our evaluation framework is designed in a systematic manner. For each  
 882 problem, we collect 8 responses from every comparison model. For a few models (i.e., DeepSeek-  
 883 V3/R1, GPT-5-mini, Gemini-2.5-pro), we reduce the number of responses to 4 due to computational  
 884 limitations and budget constraints. We adopt two evaluation metrics: *Pass@k* and *Cons@k*. For  
 885 the *Pass@k* metric, a sample is regarded as correct if at least one of the  $k$  responses generated by  
 886 the LRM matches the ground-truth answer. For the *Cons@k* metric, we apply majority voting to  
 887 derive a consensus answer for each problem, and then compute the average accuracy across the  
 888 dataset. To ensure standardized evaluation, each problem is paired with a fixed reference answer,  
 889 and correctness is assessed through exact string matching. **We enforce that LRM<sub>s</sub> output their final**  
 890 **answers in the \boxed{\{}}** **format.** For decoding hyperparameters, we follow prior work (Guo et al.,  
 891 2025; Team, 2025b), setting `temperature`, `top_p`, `min_p`, and `max_token` to 0.6, 0.95, 0, and  
 892 32768, respectively.

## 893 F DEFINITION AND EXPLANATION

### 894 F.1 TASK DEFINITION

895 (1) **Maximum Degree node (easy):** Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , the task is to  
 896 identify the node  $v \in \mathcal{V}$  with the maximum degree and output its degree  $\deg(v) = \max_{u \in \mathcal{V}} \deg(u)$ .

897 **Optimal complexity:**  $O(|V| + |E|)$  using a single scan of graph.

898 (2) **Maximum Weight Triangle (medium):** Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with a  
 899 positive node weight function  $w : \mathcal{V} \rightarrow \mathbb{R}^+$ , the task is to find a triangle  $\{v_1, v_2, v_3\}$  that maximizes  
 900 the sum of their weights  $\sum_{i=1}^3 w(v_i)$ , and output this maximum sum.

901 **Optimal complexity:**  $O(|V|^3)$  by brute-force enumeration of all triples of vertices.

902 (3) **Maximum Clique Problem (hard):** Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , the task is  
 903 to find a subset of nodes  $\mathcal{V}' \subseteq \mathcal{V}$  that forms a clique (i.e., every pair of nodes in  $\mathcal{V}'$  is connected by an  
 904 edge) with the maximum possible size  $|\mathcal{V}'|$ . The output is this size.

905 **Optimal complexity:** NP-hard. The best known exact exponential-time algorithms (e.g., branch-  
 906 and-bound, branch-and-reduce) run in about  $O(1.1996^{|V|})$ .

907 (4) **PathSum (easy):** Given an undirected tree  $\mathcal{T} = (\mathcal{V}, \mathcal{E})$  with a root  $r$  and edge weights  
 908  $w : \mathcal{E} \rightarrow \mathbb{R}$ , the task is to count the number of paths from the root  $r$  to a leaf node  $l$  such that the sum  
 909 of edge weights  $\sum_{(u,v) \in \text{path}(r,l)} w(u, v)$  exceeds a given threshold  $\tau$ .

910 **Optimal complexity:**  $O(|V|)$  using depth-first search (DFS).

911 (5) **Distance- $k$  (medium):** Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  and a specific node  $v \in \mathcal{V}$ ,  
 912 the task is to find the number of nodes  $u \in \mathcal{V}$  such that their shortest path distance  $d(u, v) \leq k$ .

918 **Optimal complexity:**  $O(|V| + |E|)$  using BFS.  
 919

920 (6) **Diameter (hard):** Given an unweighted undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , the task is to find the  
 921 longest shortest path between any two nodes  $u, v \in \mathcal{V}$  and output its length, i.e.,  $\max_{u,v \in \mathcal{V}} d(u, v)$ .  
 922

923 **Optimal complexity:**  $O(|V|(|V| + |E|))$  using BFS (or DFS) from every node.  
 924

925 (7) **Maximum  $k$ -core (easy):** Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  and an integer  $k$ , the task  
 926 is to find a subgraph  $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$  where for every node  $v \in \mathcal{V}'$ , its degree within the subgraph  
 927  $\deg_{\mathcal{G}'}(v) \geq k$ . The goal is to find the subgraph with the maximum number of nodes  $|\mathcal{V}'|$  and output  
 928 this size.  
 929

930 **Optimal complexity:**  $O(|V| + |E|)$  using the peeling algorithm.  
 931

932 (8) **Minimum Spanning Tree (medium):** Given a connected, undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$   
 933 with positive edge weights  $w : \mathcal{E} \rightarrow \mathbb{R}^+$ , the task is to find a spanning tree (a subset of edges  $\mathcal{E}' \subseteq \mathcal{E}$   
 934 that connects all nodes) with the minimum possible total edge weight  $\sum_{(u,v) \in \mathcal{E}'} w(u, v)$ . The output  
 935 is this minimum total weight.  
 936

937 **Optimal complexity:** Kruskal's algorithm:  $O(|E| \log |V|)$ .  
 938

939 (9) **Distance Threshold (hard):** Given an undirected graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  with positive edge  
 940 weights, the task is to find the node  $v^* = \arg \min_{v \in \mathcal{V}} |\{u \in \mathcal{V} \setminus \{v\} \mid d(u, v) \leq \tau\}|$ , where  $\tau$  is a  
 941 given threshold, and output the node  $v^*$ .  
 942

943 **Optimal complexity:** worst case  $O(|V||E| \log |V|)$  using Dijkstra's algorithm from each node  
 944 (truncated at distance  $\tau$ ).  
 945

## 946 F.2 ERROR TYPE CLASSIFICATION

947 We categorize model errors into four major classes:  
 948

949 1. *Algorithm Selection Error.* These errors capture whether LRM<sup>s</sup> truly understand how to  
 950 select an appropriate algorithm for a given problem.  
 951

952 (I) **Incorrect algorithm:** The chosen algorithm cannot produce the correct solution for  
 953 the problem.  
 954 (II) **Suboptimal algorithm:** The algorithm is intuition (not guaranteed to produce the  
 955 optimal solution) or computationally inefficient.  
 956

957 2. *Algorithm Execution Error.* These errors occur when LRM<sup>s</sup> make mistakes during the  
 958 execution of an algorithm.  
 959

960 (I) **State update error:** Incorrectly updating intermediate results, either by missing  
 961 necessary information or storing extraneous data.  
 962 (II) **Missing elements:** Failure to include essential steps, such as neglecting to traverse all  
 963 nodes or edges.  
 964 (III) **Condition misjudgment:** Errors in conditional checks, e.g., incorrect `if`-statements  
 965 or loop conditions.  
 966

967 3. *Output Quality Error.* In these cases, the model produces the correct answer but with  
 968 low-quality output. We primarily focus on:  
 969

970 (I) **Redundancy:** Unnecessary repetition or superfluous information in the solution.  
 971

972 4. *Information Error.* These errors arise when the model fails to capture or recall critical  
 973 problem-specific information. We focus on:  
 974

975 (I) **Graph memorization error:** Misunderstanding or incorrectly recalling the graph  
 976 structure, such as introducing extra nodes/edges or omitting existing ones.  
 977

972 F 3 DEFINITION OF ENTROPY

973

Formally, we define the entropy  $H_t$  at generation step  $t$  as the Shannon entropy of the model's predicted probability distribution over its vocabulary:

976

$$H_t = - \sum_{j=1}^V p_{t,j} \log p_{t,j} \quad (3)$$

979

Here,  $p_t = (p_{t,1}, \dots, p_{t,V})$  represents the probability distribution over the entire vocabulary  $V$ . This distribution is produced by the language model, denoted as  $\pi_\theta$ , conditioned on the input query  $q$  and the preceding token sequence  $o_{<t}$ . Specifically,  $p_t$  is computed by applying a temperature-scaled Softmax function to the model's raw output logits  $z_t$ :

983

$$p_t = \pi_\theta(\cdot \mid q, o_{<t}) = \text{Softmax}\left(\frac{z_t}{T}\right) \quad (4)$$

985

## G ADDITIONAL TASKS DESCRIPTION

988

### Problem Description G.1: Diameter

990

You are required to calculate the diameter of an undirected knowledge graph. The diameter of a graph is the greatest shortest-path distance between any two nodes in the graph.

992

## Problem to Solve

- Entities in this knowledge graph: Seine-et-Marne, Vendrest, Paris, Clickteam, Soignolles-en-Brie, Saint-Fiacre Seine-et-Marne, France, Didier Julia
- The relationships between these entities are as follows:
  - Seine-et-Marne is connected to Vendrest via the relationship department.
  - Seine-et-Marne is connected to Saint-Fiacre Seine-et-Marne via the relationship department.
  - Seine-et-Marne is connected to Didier Julia via the relationship region.
  - Seine-et-Marne is connected to Soignolles-en-Brie via the relationship department.
  - Seine-et-Marne is connected to France via the relationship country.
  - Vendrest is connected to France via the relationship country.
  - Paris is connected to France via the relationship capital.
  - Paris is connected to Clickteam via the relationship locationCity.
  - Paris is connected to Didier Julia via the relationship birthPlace.
  - Soignolles-en-Brie is connected to France via the relationship country.
  - Saint-Fiacre Seine-et-Marne is connected to France via the relationship country.
  - France is connected to Didier Julia via the relationship region.

101

Please determine the diameter of this network and output the diameter in the following format:  
`\boxed{\text{d}}`

101

### Problem Description G.2: Distance Threshold

1010

You are given an undirected weighted graph representing the airport network, where nodes represent airports (codes: DLY, TAH, HIR, IPA, VLI, BNE, NAN, AKL, LNE, FTA, AWD, EAE, LNB, AUY) and edges represent direct flights between airports. The weight of each edge is the distance between two airports.

1020

- The list of airports: DLY, TAH, HIR, IPA, VLI, BNE, NAN, AKL, LNE, FTA, AWD, EAE, LNB, AUY
- The direct flights (edges) are: [DLY, IPA, 34], [DLY, VLI, 139], [TAH, AUY, 105], [TAH, AWD, 46], [TAH, FTA, 105], [TAH, IPA, 64], [TAH, VLI, 217], [HIR, VLI, 1282], [HIR, BNE, 2126], [HIR, NAN, 2093], [IPA, VLI, 167], [VLI, NAN, 966], [VLI, AKL, 105]

1026 2238], [VLI, BNE, 1893], [VLI, EAE, 67], [VLI, LNB, 125],  
 1027 [VLI, LNE, 204], [BNE, AKL, 2295], [BNE, NAN, 2710], [NAN,  
 1028 AKL, 2156], [FTA, AWD, 72]

1029 • The distance threshold is 284.

1030 For each airport, the distance to another airport is defined as the sum of the weights (distances)  
 1031 along the shortest path connecting them.

1032 Your task is: Return the airport code with the smallest number of other airports that can be reached  
 1033 with a shortest path distance no more than the threshold. If there are multiple such airports, return  
 1034 the one with the lexicographically largest code.

1035 Present your answer in the following format: `\boxed{airport_code}`. The `airport_code` is the airport code.

### 1038 Problem Description G.3: Maximum Clique Problem

1040 You are required to solve the Maximum Clique Problem for an undirected wikipedia network. In  
 1041 this network, nodes represent wikipedia articles and edges represent hyperlinks between articles.  
 1042 Your objective is to find the largest subset of nodes such that every pair of vertices in this subset is  
 1043 connected by an edge.

1044 • Articles in the network: Eugene Domingo, Philippines, Talk show, David Cook (singer),  
 1045 Live television, BB Gandanghari, Cool Center, Katya Santos, David Archuleta, Sadako,  
 1046 Gladys Guevarra, List of Philippine television shows

1047 • Hyperlinks between these articles: Eugene Domingo and Cool Center, Eugene Domingo  
 1048 and Philippines, Philippines and Cool Center, Philippines and Gladys Guevarra, Philip-  
 1049 pines and BB Gandanghari, Philippines and List of Philippine television shows, Philip-  
 1050 pines and David Archuleta, Talk show and Cool Center, David Cook (singer) and Cool  
 1051 Center, David Cook (singer) and David Archuleta, Live television and Cool Center, BB  
 1052 Gandanghari and Cool Center, Cool Center and List of Philippine television shows, Cool  
 1053 Center and Sadako, Cool Center and David Archuleta, Cool Center and Gladys Guevarra,  
 1054 Cool Center and Katya Santos.

1055 Identify the clique with the maximum number of articles in this network. Present your answer in  
 1056 the following format: `\boxed{k}`. k is the number of articles of this clique.

### 1058 Problem Description G.4: Minimum Spanning Tree

1059 You are required to solve the Minimum Spanning Tree Problem for an undirected street network.  
 1060 In this network, nodes represent streets (e.g., street IDs) and edges represent intersections between  
 1061 streets. The weight of each edge is the distance between two streets.

1062 • Streets in the network: Mulgray Avenue, Vanessa Avenue, Eames Avenue, Gabrielle  
 1063 Avenue, Justine Avenue, Coronation Road, Turon Avenue, Jasper Road, Louise Avenue,  
 1064 Glanmire Road, Hilda Road, Seven Hills Road

1065 • Intersections between these streets: Mulgray Avenue and Coronation Road (weight: 6),  
 1066 Mulgray Avenue and Jasper Road (weight: 3), Vanessa Avenue and Jasper Road (weight:  
 1067 1), Eames Avenue and Hilda Road (weight: 7), Eames Avenue and Jasper Road (weight:  
 1068 10), Gabrielle Avenue and Justine Avenue (weight: 6), Gabrielle Avenue and Turon  
 1069 Avenue (weight: 7), Justine Avenue and Jasper Road (weight: 1), Justine Avenue and  
 1070 Turon Avenue (weight: 6), Coronation Road and Jasper Road (weight: 9), Turon Avenue  
 1071 and Jasper Road (weight: 4), Jasper Road and Glanmire Road (weight: 5), Jasper Road  
 1072 and Hilda Road (weight: 1), Jasper Road and Louise Avenue (weight: 5), Jasper Road  
 1073 and Seven Hills Road (weight: 2), Hilda Road and Seven Hills Road (weight: 1).

1074 Identify the minimum spanning tree of this network. The minimum spanning tree is a subset of  
 1075 edges in a connected, weighted graph that connects all the vertices together with the smallest  
 1076 possible total edge weight and without any cycles.

1077 Present your answer in the following format: `\boxed{n}`, where n is the sum of the weights of  
 1078 the edges in the minimum spanning tree.

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1097

1098

1099

1100

1101

1102

1103

1104

1105

1106

1107

**Problem Description G.5: PathSum**

You are given a binary tree representing a co-authorship network. Each node is an author, and each edge represents a co-authorship relationship, with the edge's weight indicating the number of papers co-authored by the two authors. Find all paths from the root author to any leaf author such that the sum of the edge weights (i.e., the total number of co-authored papers along the path) is greater than the given value.

- Authors in the network: Tarkan Tan, Ruud H. Teunter, Hui-Ming Wee, Samuel Yáñez Artus, Po-Chung Yang, Asoke Kumar Bhunia, Anne Barros, Yu-Chung Tsao, Samiran Chattopadhyay, A. K. Bhunia, Michel Roussignol, Mahmood Shafiee, Shib Sankar Sana, Gwo-Ji Sheen
- Co-authorship relationships (with number of co-authored papers): Tarkan Tan and Ruud H. Teunter (co-authored 2 papers), Ruud H. Teunter and Hui-Ming Wee (co-authored 4 papers), Ruud H. Teunter and Samuel Yáñez Artus (co-authored 9 papers), Hui-Ming Wee and Po-Chung Yang (co-authored 4 papers), Hui-Ming Wee and Asoke Kumar Bhunia (co-authored 5 papers), Samuel Yáñez Artus and Anne Barros (co-authored 6 papers), Po-Chung Yang and Yu-Chung Tsao (co-authored 4 papers), Asoke Kumar Bhunia and Samiran Chattopadhyay (co-authored 6 papers), Asoke Kumar Bhunia and A. K. Bhunia (co-authored 1 papers), Anne Barros and Michel Roussignol (co-authored 9 papers), Anne Barros and Mahmood Shafiee (co-authored 5 papers), Yu-Chung Tsao and Shib Sankar Sana (co-authored 10 papers), Yu-Chung Tsao and Gwo-Ji Sheen (co-authored 3 papers).
- The root of the tree is Tarkan Tan.
- The target value is 19.

Present your answer in the following format: `\boxed{n}`. n is the number of qualifying paths.

**Problem Description G.6: Maximum Weight Triangle**

You are required to solve the Maximum Triangle Sum Problem for an undirected wikipedia network. In this network, nodes represent wikipedia articles and edges represent hyperlinks between articles. Each node is assigned a weight. Your objective is to find the triangle with the maximum sum of weights of its three nodes.

- Articles in the network: Montebelluna (weight: 1), Massimo Mascioletti (weight: 8), Gianluca Faliva (weight: 4), Eppelheim (weight: 7), United States (weight: 6), Alberto Rebecca (weight: 10), October 15 (weight: 6), Rugby union (weight: 1), Italy (weight: 9), List of football clubs in Italy (weight: 9), Manuel Dallan (weight: 8), Brad Johnstone (weight: 1)
- Hyperlinks between these articles: Montebelluna and List of football clubs in Italy, Montebelluna and Eppelheim, Montebelluna and Alberto Rebecca, Montebelluna and Italy, Montebelluna and Manuel Dallan, Massimo Mascioletti and Brad Johnstone, Massimo Mascioletti and Gianluca Faliva, Gianluca Faliva and Italy, Gianluca Faliva and Brad Johnstone, Gianluca Faliva and Manuel Dallan, Gianluca Faliva and Rugby union, Eppelheim and Italy, United States and Italy, United States and October 15, Alberto Rebecca and Italy, October 15 and Italy, October 15 and Manuel Dallan, Rugby union and Manuel Dallan, Italy and Manuel Dallan.

Identify the triangle with the maximum sum of weights of its three nodes in this network. Present your answer in the following format: `\boxed{n}`. n is the maximum sum of weights of its three nodes.

**Problem Description G.7: Minimum Spanning Tree (Adjacency List)**

You are required to solve the Minimum Spanning Tree Problem for an undirected street network. In this network, nodes represent streets (e.g., street IDs) and edges represent intersections between streets. The weight of each edge is the distance between two streets.

1134

- Streets in the network: Maplin Road, Golden Plover Close, Widgeon Close, Downs Court
- 1135 Road, Greyfields Close, Pheasant Close, Coolfin Road, Burrard Road, Freemasons Road,
- 1136 Ethel Road, Partridge Knoll
- 1137
- Adjacency list representation (each street followed by its connected streets): Maplin Road
- 1138 (weight: 5): Freemasons Road (weight: 9), Golden Plover Close (weight: 6), Pheasant
- 1139 Close (weight: 8), Widgeon Close (weight: 7)
- 1140 Golden Plover Close (weight: 4): Maplin Road (weight: 5), Widgeon Close (weight: 7)
- 1141 Widgeon Close (weight: 7): Maplin Road (weight: 5), Golden Plover Close (weight: 4)
- 1142 Downs Court Road (weight: 3): Partridge Knoll (weight: 10)
- 1143 Greyfields Close (weight: 6): Partridge Knoll (weight: 10)
- 1144 Pheasant Close (weight: 8): Maplin Road (weight: 5), Partridge Knoll (weight: 10)
- 1145 Coolfin Road (weight: 2): Freemasons Road (weight: 9)
- 1146 Burrard Road (weight: 4): Freemasons Road (weight: 9)
- 1147 Freemasons Road (weight: 9): Maplin Road (weight: 5), Coolfin Road (weight: 2),
- 1148 Burrard Road (weight: 4), Ethel Road (weight: 3) Ethel Road (weight: 3): Freemasons
- 1149 Road (weight: 9)
- 1150 Partridge Knoll (weight: 10): Downs Court Road (weight: 3), Greyfields Close (weight:
- 1151 6), Pheasant Close (weight: 8)

1152

1153 Identify the minimum spanning tree of this network. The minimum spanning tree is a subset of

1154 edges in a connected, weighted graph that connects all the vertices together with the smallest

1155 possible total edge weight and without any cycles.

1156 Present your answer in the following format:  $\boxed{n}$ , where  $n$  is the sum of the weights of

1157 the edges in the minimum spanning tree.

1158

1159

1160

#### 1161 Problem Description G.8: Minimum Spanning Tree (Markdown Table)

1162

1163

1164

1165

1166

1167   Street 1	1168   Street 2	1169   Weight
1170   Heritage Court	1171   Darcey Road	1172   3
1173   Jordana Place	1174   Darcey Road	1175   6
1176   Neville Court	1177   Darcey Road	1178   7
1179   Silky Oak Place	1180   Darcey Road	1181   3
1182   Candlebush Crescent	1183   Darcey Road	1184   1
1185   Candlebush Crescent	1186   Henley Close	1187   4
1188   Candlebush Crescent	1189   Lemonwood Place	1190   10
1191   Darcey Road	1192   Melaleuca Close	1193   9
1194   Darcey Road	1195   Castlewood Drive	1196   10
1197   Darcey Road	1198   Crane Road	1199   4
1200   Darcey Road	1201   Henley Close	1202   5
1203   Darcey Road	1204   Jarrah Place	1205   10
1206   Crane Road	1207   Castlewood Drive	1208   8

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223 Find the minimum spanning tree of this street network and output the sum of the weights of the

1224 edges in the MST. Present your answer in the following format:  $\boxed{n}$ , where  $n$  is the sum of the

1225 weights of the edges in the MST.

1188

**Problem Description G.9: Minimum Spanning Tree (Numeric ID)**

1189

You are given an undirected weighted graph with 12 nodes labeled from 0 to 11.

1190

- Edges (format: [Node1, Node2, weight]): [0, 5, 3], [1, 5, 6], [2, 5, 7], [3, 5, 3], [4, 5, 1], [4, 6, 4], [4, 7, 10], [4, 8, 9], [5, 11, 10], [5, 10, 4], [5, 6, 5], [5, 9, 10], [10, 11, 8]

1191

Find the minimum spanning tree of this graph and output the sum of the weights of the edges in the MST. Present your answer in the following format: ' $\boxed{n}$ ', where n is the sum of the weights of the edges in the minimum spanning tree.

1192

**Problem Description G.10: Minimum Spanning Tree (Code Version)**

1193

You are required to solve the Maximum Triangle Sum Problem for an undirected wikipedia network. In this network, nodes represent wikipedia articles and edges represent hyperlinks between articles. Each node is assigned a weight. Your objective is to find the triangle with the maximum sum of weights of its three nodes.

1194

- Streets in the network: Mulgray Avenue, Vanessa Avenue, Eames Avenue, Gabrielle Avenue, Justine Avenue, Coronation Road, Turon Avenue, Jasper Road, Louise Avenue, Glammire Road, Hilda Road, Seven Hills Road
- Intersections between these streets: Mulgray Avenue and Coronation Road (weight: 6), Mulgray Avenue and Jasper Road (weight: 3), Vanessa Avenue and Jasper Road (weight: 1), Eames Avenue and Hilda Road (weight: 7), Eames Avenue and Jasper Road (weight: 10), Gabrielle Avenue and Justine Avenue (weight: 6), Gabrielle Avenue and Turon Avenue (weight: 7), Justine Avenue and Jasper Road (weight: 1), Justine Avenue and Turon Avenue (weight: 6), Coronation Road and Jasper Road (weight: 9), Turon Avenue and Jasper Road (weight: 4), Jasper Road and Glammire Road (weight: 5), Jasper Road and Hilda Road (weight: 1), Jasper Road and Louise Avenue (weight: 5), Jasper Road and Seven Hills Road (weight: 2), Hilda Road and Seven Hills Road (weight: 1).

1195

Identify the minimum spanning tree of this network. The minimum spanning tree is a subset of edges in a connected, weighted graph that connects all the vertices together with the smallest possible total edge weight and without any cycles.

1196

Solve this problem using Python code only. Wrap your Python code in the following format:

1197

```
# Your code here
print ("\\boxed{n}")
```

1198

Here, 'n' is the sum of the weights of the edges in the minimum spanning tree.

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

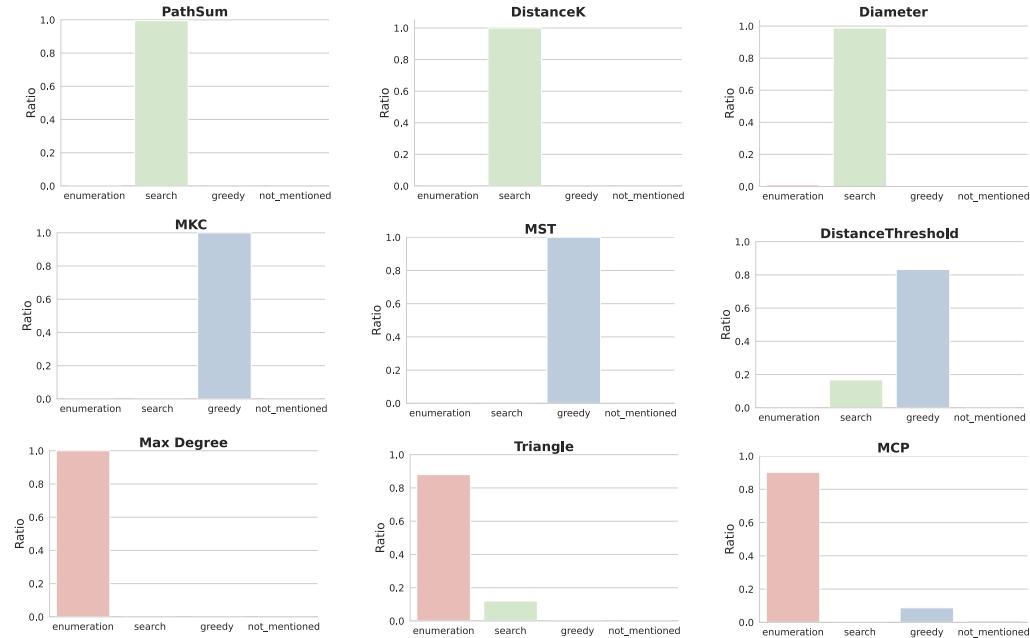
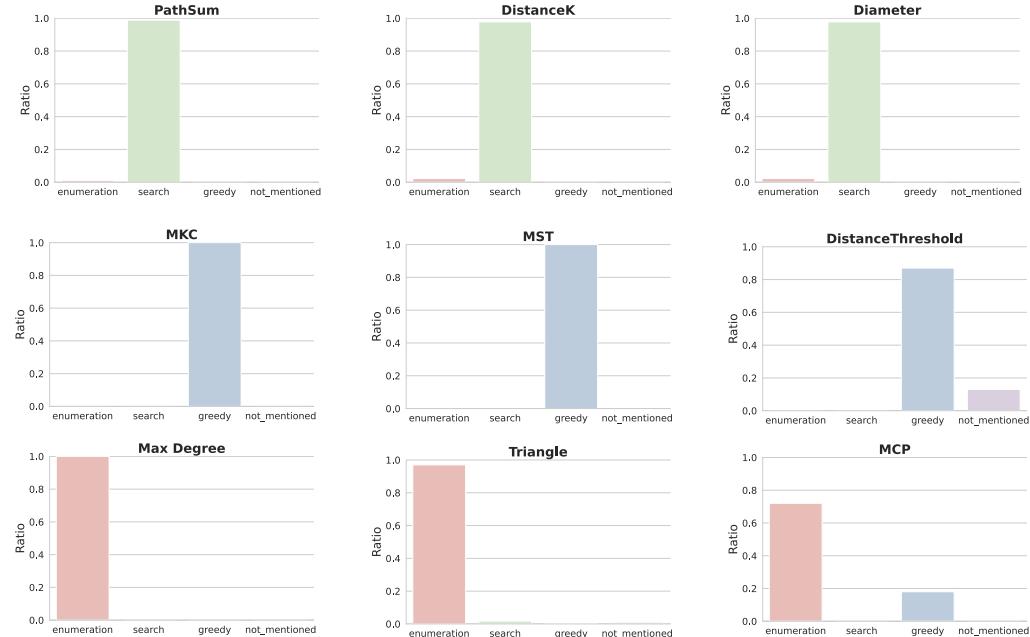
1239

1240

1241

## 1242 H ADDITIONAL EXERIMENTS RESULTS

### 1244 H.1 ALGORITHM RATIO CALCULATION

1267 Figure 8: Algorithm ratio of Qwen3-32B.  
12681291 Figure 9: Algorithm ratio of Gemini-2.5-Pro.  
12921293  
1294  
1295

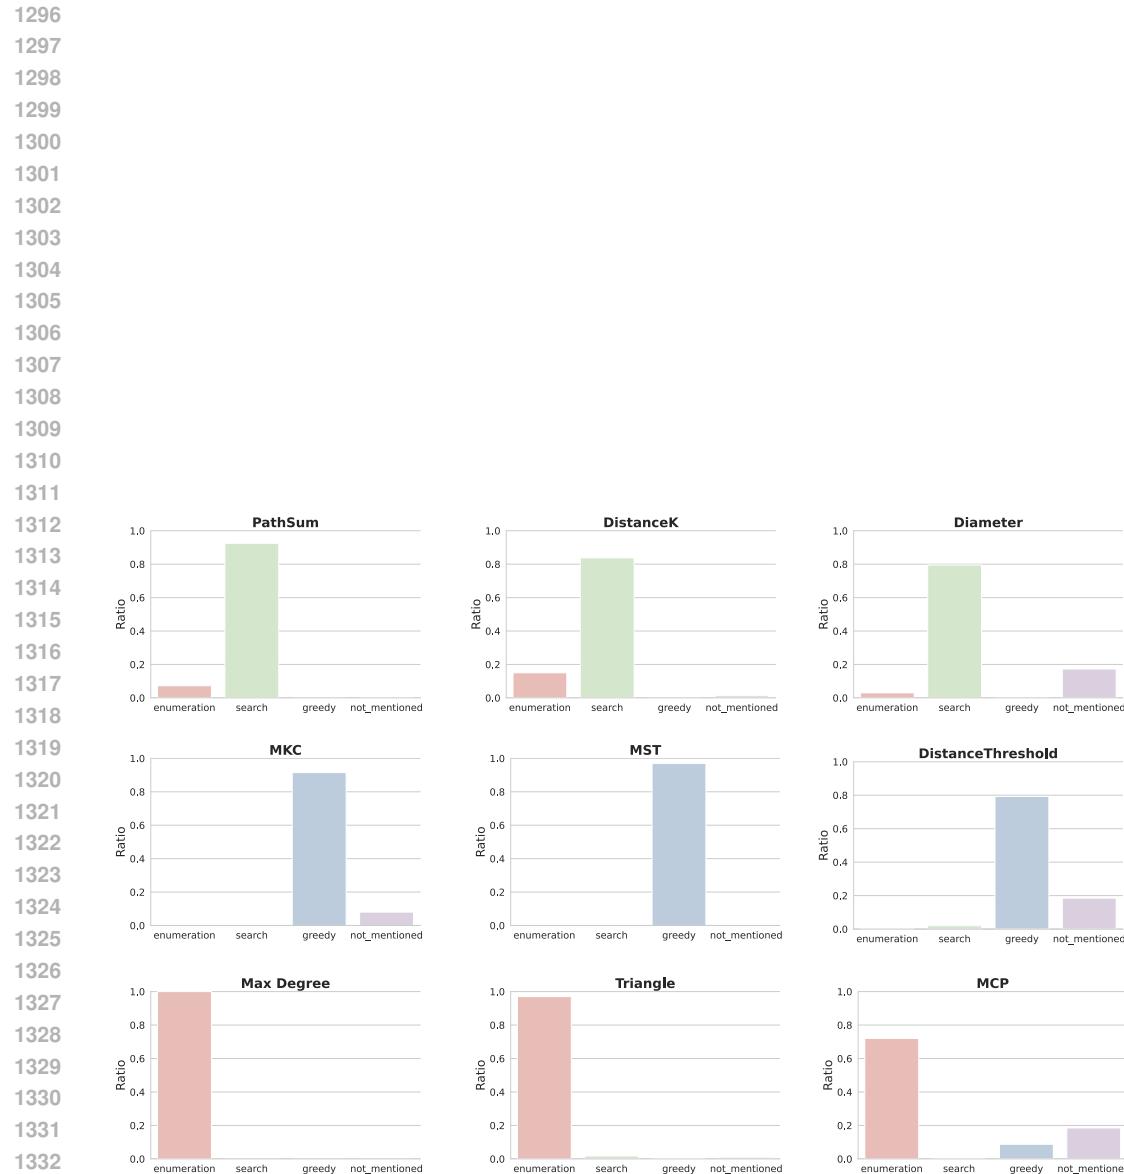


Figure 10: Algorithm ratio of Llama-3.3-70B.

1350  
1351 H.2 PERFORMANCE ACROSS GRAPH ALGORITHM CATEGORIES (RQ1)  
1352  
1353

Table 5: Evaluation results of Level-1 and Level-2 groups.

Models	Level-1 (8-15 nodes)								Level-2 (16-30 nodes)							
	Enumeration		Exploration		Intuition		Avg.		Enumeration		Exploration		Intuition		Avg.	
	c@k	p@k	c@k	p@k	c@k	p@k	c@k	p@k	c@k	p@k	c@k	p@k	c@k	p@k	c@k	p@k
Skywork-OR1-7B-Preview	0.50	0.91	0.63	0.93	0.31	0.67	0.48	0.84	0.20	0.67	0.20	0.67	0.08	0.39	0.16	0.58
Light-R1-7B-DS	0.37	0.79	0.41	0.82	0.24	0.61	0.34	0.74	0.07	0.48	0.09	0.49	0.04	0.23	0.07	0.40
Distill-Qwen-7B	0.43	0.88	0.47	0.87	0.30	0.68	0.40	0.81	0.17	0.68	0.14	0.63	0.02	0.32	0.11	0.54
Qwen2.5-7B	0.21	0.47	0.29	0.65	0.06	0.20	0.19	0.44	0.12	0.45	0.13	0.29	0.04	0.13	0.10	0.29
OpenThinker-7B	0.57	0.88	0.54	0.91	0.41	0.79	0.51	0.86	0.23	0.83	0.17	0.61	0.10	0.37	0.17	0.60
Qwen3-8B-no-thinking	0.57	0.85	0.70	0.96	0.30	0.65	0.52	0.82	0.22	0.62	0.35	0.75	0.03	0.29	0.20	0.55
Qwen3-8B	0.89	0.99	0.77	0.99	0.96	1.00	0.87	0.99	0.63	0.93	0.63	0.93	0.66	0.93	<b>0.64</b>	<b>0.93</b>
Light-R1-32B	0.97	1.00	0.97	1.00	0.95	0.99	0.96	<b>1.00</b>	0.85	0.99	0.89	0.97	0.71	0.93	0.82	0.96
Skywork-OR1-32B	0.91	0.99	1.00	1.00	0.89	1.00	0.94	<b>1.00</b>	0.93	0.99	0.97	0.99	0.87	0.98	<b>0.92</b>	<b>0.99</b>
Distill-Qwen-32B	0.93	0.99	0.96	0.99	0.96	1.00	0.95	0.99	0.75	0.98	0.73	0.95	0.62	0.89	0.70	0.94
Qwen2.5-32B	0.25	0.73	0.42	0.85	0.36	0.61	0.34	0.73	0.11	0.56	0.18	0.52	0.16	0.39	0.15	0.49
OpenThinker-32B	0.97	0.99	0.97	1.00	0.95	1.00	0.96	<b>1.00</b>	0.85	0.98	0.85	0.97	0.62	0.94	0.77	0.96
QWQ-32B	0.98	1.00	1.00	1.00	0.99	1.00	<b>0.99</b>	<b>1.00</b>	0.93	0.99	0.92	0.99	0.91	0.98	<b>0.92</b>	<b>0.99</b>
Qwen3-32B	0.97	0.99	0.98	1.00	0.99	1.00	<b>0.98</b>	<b>1.00</b>	0.89	0.99	0.83	0.98	0.84	0.96	0.85	0.98
Qwen3-32B-no-thinking	0.77	0.97	0.84	0.97	0.66	0.87	0.76	0.94	0.51	0.84	0.59	0.89	0.19	0.51	0.43	0.75

1363  
1364  
1365  
1366  
1367  
1368  
1369  
1370 **Normalized Analysis.** To address the imbalance in task difficulty across reasoning taxonomies, we  
1371 follow prior work (Yuan et al., 2024; Yu et al., 2023) and apply a normalization procedure to reduce  
1372 the influence of task complexity on model performance. Specifically, for each model and task, we  
1373 compute a  $z$ -score of Pass@ $k$  and Cons@ $k$  accuracy and then linearly rescale the scores to the range  
1374 [0, 100]. Specifically, for each model  $i$  and each task  $j$ , we compute the  $z$ -score of its Pass@ $k$  and  
1375 Cons@ $k$  accuracy as follows:

$$z_{ij} = \frac{x_{ij} - \mu(x_{i1}, \dots, x_{i|M|})}{\sigma(x_{i1}, \dots, x_{i|M|})}, \quad (5)$$

1376 where  $x_{ij}$  denotes the Pass@ $k$  or Cons@ $k$  accuracy of model  $i$  on task  $j$ , while  $\mu(\cdot)$  and  $\sigma(\cdot)$  denote  
1377 the mean and standard deviation over all  $M$  models on that task.

1378 We then linearly rescale the  $z$ -scores into the range [0, 100] to obtain a normalized score:

$$s_{ij} = 100 \cdot \frac{z_{ij} - \min(z)}{\max(z) - \min(z)}, \quad (6)$$

1379 where  $\min(z)$  and  $\max(z)$  are taken over all models on the same task.

1380 Figure 11-Figure 14 present the normalized scores of different models across reasoning taxonomies  
1381 and graph sizes. The normalized outcomes preserve the same performance hierarchy observed earlier,  
1382 highlighting persistent weaknesses of LRM in intuitive reasoning.

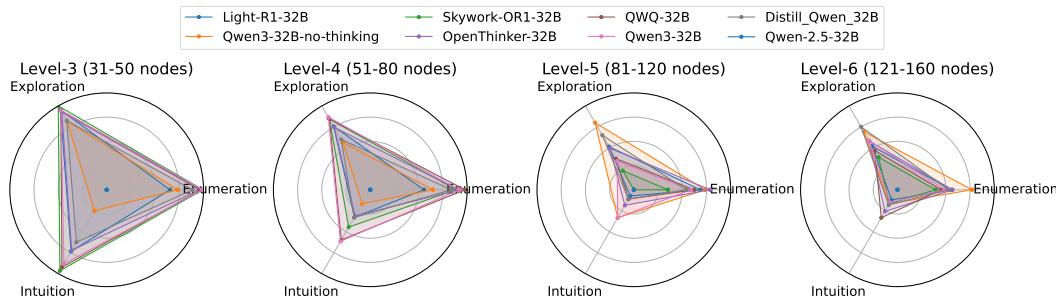


Figure 11: Normalized pass@k scores of different reasoning taxonomies (32B models).

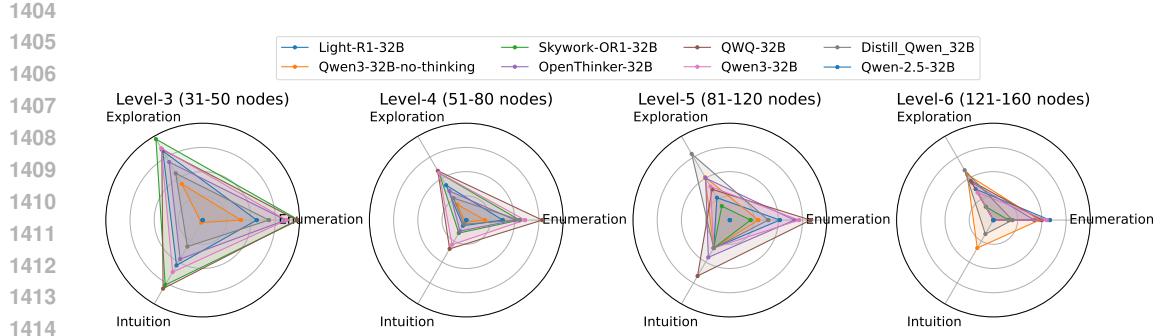


Figure 12: Normalized cons@k z-scores of different reasoning taxonomies (32B models).

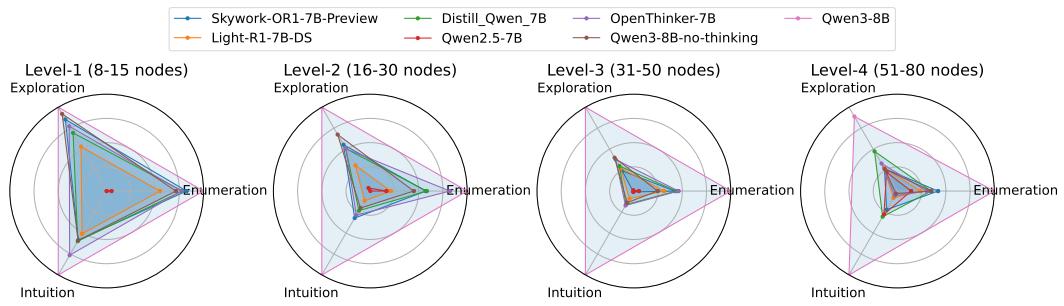


Figure 13: Normalized pass@k z-scores of different reasoning taxonomies (8B models).

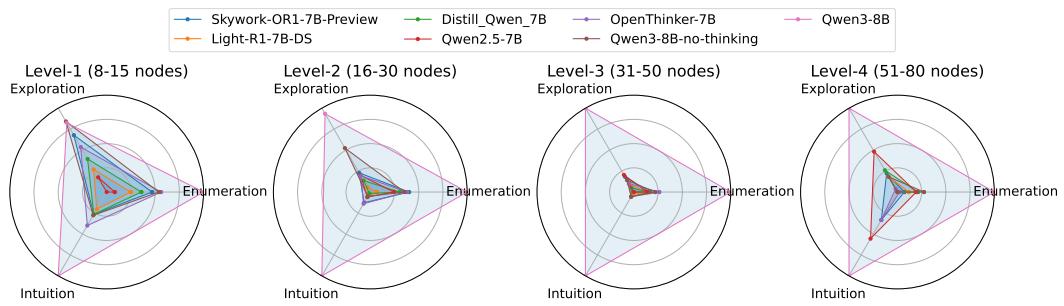


Figure 14: Normalized cons@k Z-scores of different reasoning taxonomies (8B models).

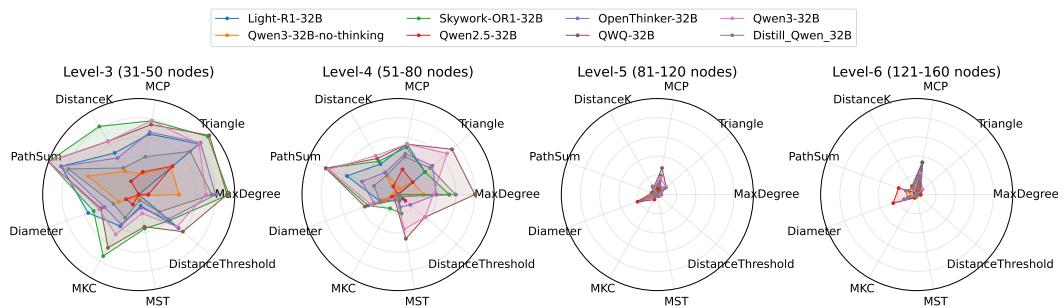


Figure 15: Cons@k accuracy of different reasoning taxonomies (32B models).

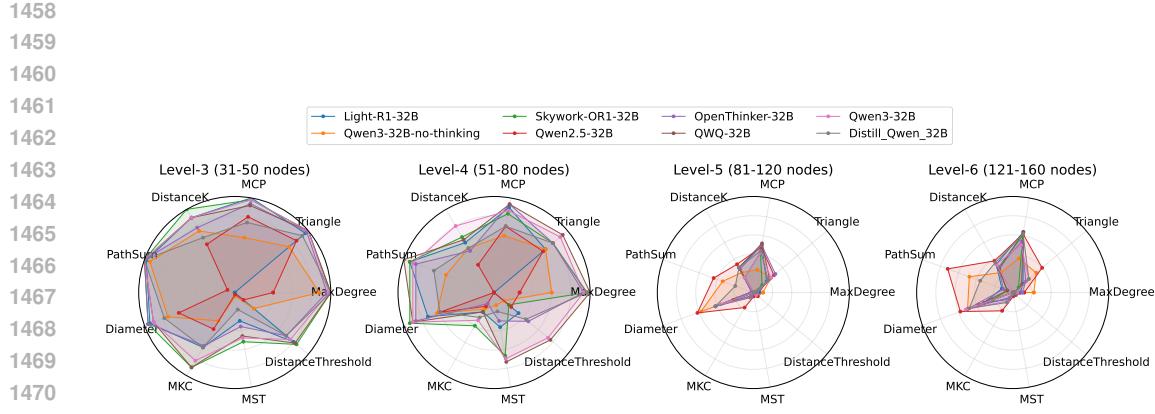


Figure 16: Pass@k accuracy of different reasoning taxonomies (32B models).

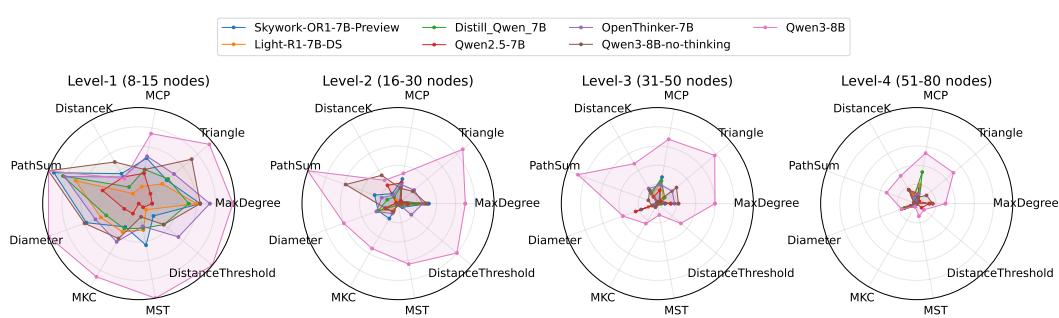


Figure 17: Cons@k accuracy of different reasoning taxonomies (8B models).

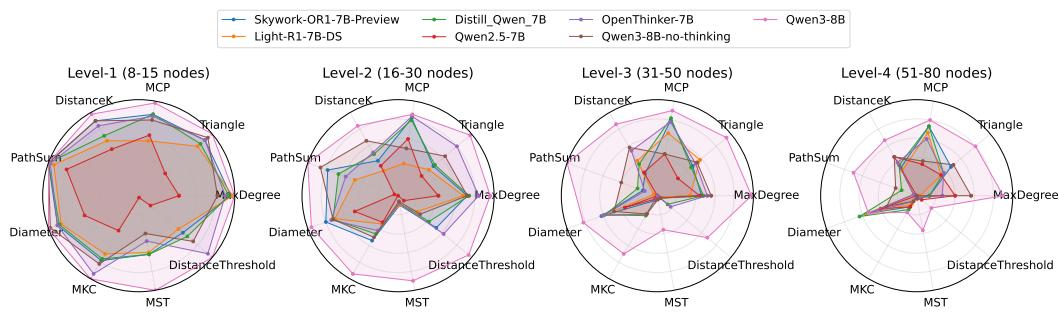
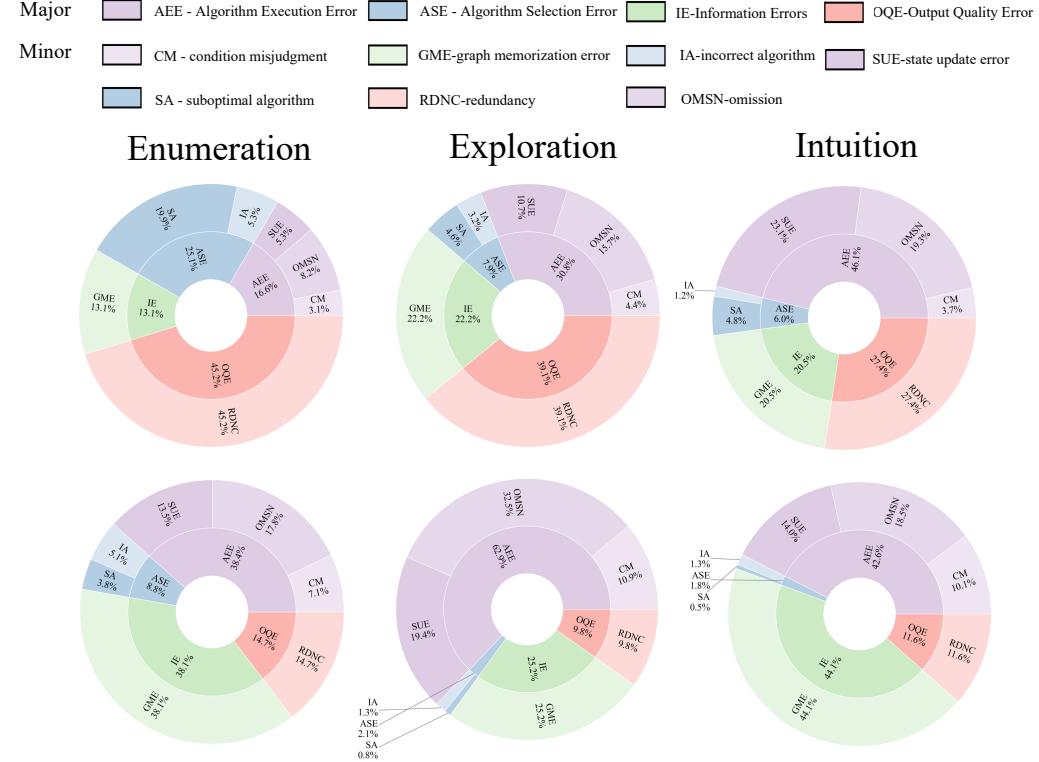
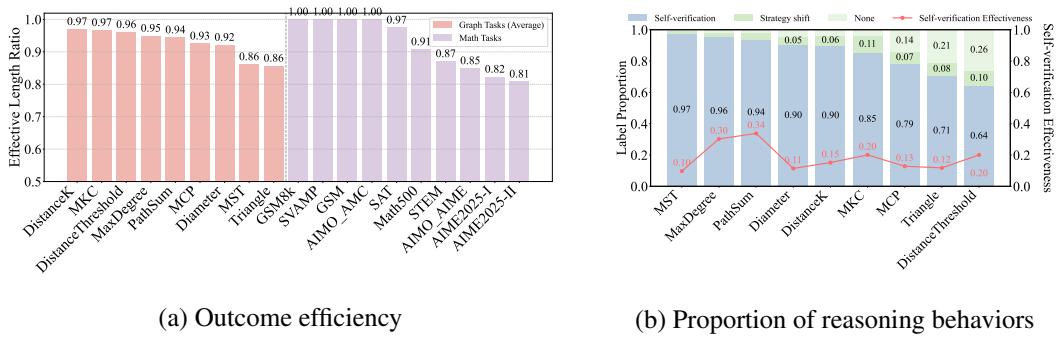


Figure 18: Pass@k accuracy of different reasoning taxonomies (8B models).

1512 H.3 ERROR ANALYSIS (RQ1)  
15131540 Figure 19: Error type distributions across reasoning taxonomies for Qwen3-8B (top)  
1541 and its non-reasoning variant Qwen3-8B-no-thinking (bottom).1542 H.4 OVERTHINKING ANALYSIS (RQ2)  
15431557 Figure 20: Overthinking analysis of Distill-Qwen-32B.  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565

1566

1567

1568

1569

1570

1571

1572

1573

1574

1575

1576

1577

1578

1579

1580

1581

1582

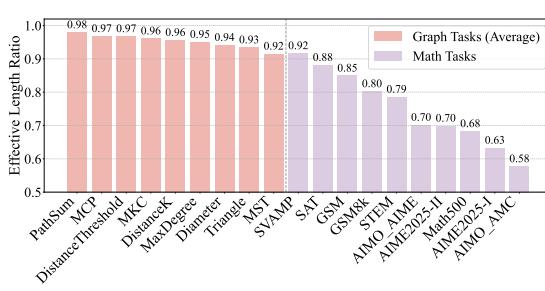
1583

1584

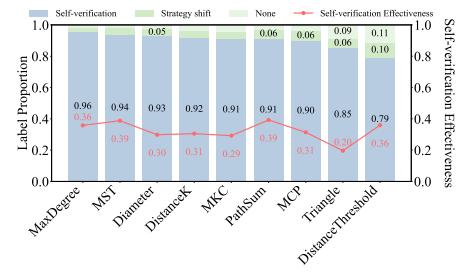
1585

1586

1587



(a) Outcome efficiency



(b) Proportion of reasoning behaviors

Figure 21: Overthinking analysis of Qwen3-32B.

1597

1598

1599

1600

1601

1602

1603

1604

1605

1606

1607

1608

1609

1610

1611

1612

1613

1614

1615

1616

1617

1618

1619

1620 

## I DISCUSSION

1621

1622 Table 6: Performance (pass@k / cons@k (%)) of Qwen3-32B under different graph representations.
1623

1624

Level	Task	Edge list	Adjacency list	Markdown table	Numeric IDs
Level-1 (8–15 nodes)	DistanceK	100.0 / 100.0	100.0 / 100.0	100.0 / 100.0	100.0 / 100.0
	MST	100.0 / 100.0	100.0 / 100.0	100.0 / 100.0	100.0 / 100.0
Level-2 (16–30 nodes)	DistanceK	100.0 / 83.3	100.0 / 96.7	96.7 / 86.7	100.0 / 90.0
	MST	96.7 / 83.3	100.0 / 86.7	93.3 / 83.3	96.7 / 83.3
Level-3 (31–50 nodes)	DistanceK	93.3 / 66.7	100.0 / 100.0	86.7 / 76.7	96.7 / 83.3
	MST	23.3 / 13.3	23.3 / 3.3	53.3 / 13.3	53.3 / 26.7

1632 

## J OTHER EXPERIMENTS

1633

1634 Table 7: Level-3 (31–50 nodes) results. Each cell reports pass@k / cons@k.
1635

1636

Model	Easy			Medium			Hard		
	Max Degree	Path Sum	MKC	Triangle	Distance-K	MST	MCP	Diameter	Distance Threshold
Distill_Qwen_32B	1.00 / 0.80	1.00 / 0.74	0.60 / 0.28	0.90 / 0.66	0.66 / 0.34	0.20 / 0.00	0.72 / 0.36	0.72 / 0.32	0.24 / 0.24
Light-R1-32B	0.90 / 0.84	0.96 / 0.90	0.62 / 0.34	0.94 / 0.80	0.84 / 0.50	0.33 / 0.07	1.00 / 0.66	0.92 / 0.56	0.56 / 0.50
Llama-3.3-70B	0.78 / 0.32	0.56 / 0.28	0.06 / 0.02	0.12 / 0.00	0.34 / 0.04	0.00 / 0.00	0.08 / 0.00	0.20 / 0.02	0.00 / 0.00
OpenThinker-32B	0.96 / 0.78	0.98 / 0.80	0.62 / 0.36	0.98 / 0.80	0.76 / 0.44	0.30 / 0.10	0.94 / 0.68	0.90 / 0.40	0.40 / 0.30
QWQ-32B	0.98 / 0.96	1.00 / 0.98	0.90 / 0.64	1.00 / 0.94	0.88 / 0.64	0.70 / 0.50	0.92 / 0.76	0.90 / 0.48	0.80 / 0.60
Qwen2.5-32B	0.40 / 0.12	0.08 / 0.00	0.40 / 0.12	0.82 / 0.44	0.58 / 0.18	0.00 / 0.00	0.80 / 0.26	0.58 / 0.18	0.12 / 0.00
Qwen3-32B	0.98 / 0.76	1.00 / 0.98	0.78 / 0.46	1.00 / 0.80	0.88 / 0.66	0.63 / 0.30	0.98 / 0.76	0.90 / 0.48	0.64 / 0.56
Qwen3-32B-no-thinking	0.84 / 0.36	0.94 / 0.58	0.30 / 0.02	0.68 / 0.44	0.72 / 0.32	0.13 / 0.03	0.54 / 0.20	0.66 / 0.20	0.24 / 0.02
Skywork-ORI-32B	1.00 / 0.94	1.00 / 1.00	0.90 / 0.76	1.00 / 0.92	1.00 / 0.82	0.67 / 0.27	0.98 / 0.82	0.94 / 0.52	0.54 / 0.38
gpt_lose_20b	0.98 / 0.90	1.00 / 1.00	0.82 / 0.60	0.98 / 0.88	0.94 / 0.76	0.73 / 0.50	0.96 / 0.78	0.88 / 0.44	0.42 / 0.22

1645 Table 8: Level-4 (51–80 nodes) results. Each cell reports pass@k / cons@k.
1646

1647

Model	Easy			Medium			Hard		
	Max Degree	Path Sum	MKC	Triangle	Distance-K	MST	MCP	Diameter	Distance Threshold
Distill_Qwen_32B	0.93 / 0.57	0.67 / 0.30	0.27 / 0.07	0.80 / 0.43	0.40 / 0.20	0.16 / 0.06	0.70 / 0.33	0.60 / 0.20	0.03 / 0.03
Light-R1-32B	0.60 / 0.50	0.90 / 0.53	0.23 / 0.03	0.70 / 0.37	0.57 / 0.33	0.30 / 0.12	0.87 / 0.50	0.67 / 0.23	0.07 / 0.07
Llama-3.3-70B	0.47 / 0.13	0.10 / 0.00	0.03 / 0.00	0.13 / 0.00	0.17 / 0.00	0.00 / 0.00	0.10 / 0.00	0.23 / 0.00	0.00 / 0.00
OpenThinker-32B	0.90 / 0.37	0.80 / 0.37	0.13 / 0.03	0.80 / 0.50	0.50 / 0.23	0.36 / 0.14	0.80 / 0.40	0.87 / 0.27	0.13 / 0.13
QWQ-32B	1.00 / 0.73	1.00 / 0.90	0.23 / 0.10	0.90 / 0.70	0.63 / 0.43	0.46 / 0.36	0.90 / 0.57	0.87 / 0.37	0.70 / 0.33
Qwen2.5-32B	0.27 / 0.00	0.00 / 0.00	0.17 / 0.03	0.57 / 0.20	0.30 / 0.07	0.02 / 0.00	0.70 / 0.20	0.57 / 0.07	0.20 / 0.10
Qwen3-32B	0.93 / 0.40	0.90 / 0.77	0.30 / 0.03	0.87 / 0.63	0.77 / 0.47	0.48 / 0.14	0.83 / 0.53	0.90 / 0.23	0.50 / 0.33
Qwen3-32B-no-thinking	0.60 / 0.23	0.50 / 0.07	0.13 / 0.00	0.60 / 0.17	0.53 / 0.27	0.04 / 0.00	0.53 / 0.03	0.60 / 0.23	0.13 / 0.03
Skywork-ORI-32B	0.93 / 0.53	0.93 / 0.83	0.37 / 0.17	0.80 / 0.37	0.67 / 0.37	0.52 / 0.40	0.80 / 0.50	0.93 / 0.30	0.00 / 0.00
gpt_lose_20b	0.97 / 0.83	0.80 / 0.67	0.37 / 0.17	0.70 / 0.50	0.57 / 0.47	0.34 / 0.20	0.67 / 0.43	0.63 / 0.37	0.63 / 0.50

1648 

**Findings of group problems by computational complexity** We re-evaluate model performance by grouping tasks based on theoretical time complexity—Easy ( $O(n)$ ), Medium ( $O(n^2)$ ), and Hard ( $O(n^3)$ /NP-hard)—while controlling for graph size (using Level-3 and Level-4 datasets) to isolate the impact of algorithmic difficulty. Results are shown in Table 7 and 8. This analysis yields three critical insights. First, within the same reasoning taxonomy, performance generally declines as computational complexity increases; for instance, models like QWQ-32B achieve near-perfect scores on linear-time tasks (e.g., *Max Degree*) but degrade significantly on NP-hard tasks (e.g., *Distance Threshold*), confirming that expanded search spaces pose greater challenges to LRM<sub>s</sub>. Second, we observe an anomaly in the *Exploration* category where models frequently outperform on the “Hard” *Diameter* task compared to the “Medium” *Distance-K* task (e.g., QWQ-32B scores 0.87 vs. 0.63 on Level-4), a discrepancy we attribute to training data distributions favoring specific classic problems. Third, reasoning taxonomy often outweighs theoretical complexity: models consistently perform better on “Hard” *Enumeration* tasks than “Easy” *Intuition* tasks (e.g., Light-R1-32B achieves 1.00 on the NP-hard *MCP* but only 0.62 on the linear-time *MKC*). This suggests that current LRM<sub>s</sub> are inherently optimized for systematic, step-by-step verification rather than the global heuristic judgments required for intuition-based problems, even when the theoretical complexity suggests the former should be more difficult.

1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673

1674 **K DETAILS DEMONSTRATION**  
16751676 **K.1 PROMPTS**  
16771678 **Prompts K.1: LLM Prompts For Categorizing Enumeration Algorithm**  
16791680 **SYSTEM PROMPT**1681 You are a graph theory expert. Given a graph problem and an answer to that problem, please  
1682 determine which of the following approaches is used in the answer:1683   a). Enumeration (such as brute-force methods)  
1684   b). Search (such as BFS, DFS)  
1685   c). Greedy  
1686   d). Not mentioned (other algorithmic approaches)1687 Directly output your choice (a, b, c, or d) with no explanation or additional text.  
16881689 **USER PROMPT**1690 Given the problem and its answer below, identify which of the following approaches was used to  
1691 derive the solution:

1692 Problem: {question}

1693 Answer: {answer}

1694 **Prompts K.2: LLM Prompts For Categorizing Exploration Algorithm**1695 **SYSTEM PROMPT**1696 You are a graph theory expert. Given a problem related to finding the shortest path and its  
1697 corresponding answer, please determine which of the following approaches is used in the answer:1698   a). Enumeration (such as brute-force methods)  
1699   b). Search (such as BFS, DFS)  
1700   c). Greedy (such as Dijkstra's algorithm)  
1701   d). Not mentioned (other algorithmic approaches)1702 Directly output your choice (a, b, c or d) with no explanation or additional text.  
17031704 **USER PROMPT**1705 Given the problem and its answer below, identify which of the following approaches was used to  
1706 derive the solution:

1707 Problem: {question}

1708 Answer: {answer}

1709 **Prompts K.3: LLM Prompts For Categorizing Intuition Algorithm**1710 **SYSTEM PROMPT**1711 You are a graph theory expert. Given a Maximum K-core (MKC) problem and an answer to that  
1712 problem, please determine which of the following approaches is used in the answer:1713   a). Enumeration (such as brute-force methods)  
1714   b). Search (such as BFS, DFS)  
1715   c). Greedy (such as Peeling algorithm)  
1716   d). Not mentioned (other algorithmic approaches)1717 Directly output your choice (a, b, c or d) with no explanation or additional text.  
17181719 **USER PROMPT**1720 Given the problem and its answer below, identify which of the following approaches was used to  
1721 derive the solution:

1722 Problem: {question}

1723 Answer: {answer}

1728 **Prompts K.4: LLM Prompts for Categorizing Errors In AI-Generated Solutions**  
17291730 You are an intelligent AI assistant. Given a graph problem and an LLM's response to that problem,  
1731 analyze the LLM's response to identify any errors it contains. Use the following refined error  
1732 categories and definitions for your analysis:1733 **OUTPUT QUALITY**1734 • **redundancy:** Unnecessary repetition or superfluous information in the solution.  
17351736 **ALGORITHM SELECTION**1737 • **incorrect algorithm:** The chosen algorithm cannot produce the cor-  
1738 rect solution for the problem. In this task, the correct algorithms are:  
1739 {correct\_algorithm\_dict[args.task]}.  
1740 • **suboptimal algorithm:** The algorithm used is inefficient or is imple-  
1741 mented in a less efficient way. In this task, the efficient algorithms are:  
1742 {efficient\_algorithm\_dict[args.task]}.  
17431744 **INFORMATION ERRORS**1745 • **graph memorization error:** Misunderstanding or incorrect memory of the graph struc-  
1746 ture, such as extra or missing nodes/edges.1747 **ALGORITHM EXECUTION**1748 • **state update error:** During algorithm execution, state variables or data structures  
1749 are updated incorrectly, causing subsequent steps to operate on erroneous states and  
1750 compromising overall correctness.  
1751 • **omission:** Missing important elements during execution, e.g., failure to traverse all nodes  
1752 or edges.  
1753 • **condition misjudgment:** Mistakes in condition checks, such as incorrect if-statement or  
1754 loop-condition evaluations.  
17551756 **Instructions:** The LLM's response will be segmented into sections labeled as "[Section 1],  
1757 content...", "[Section 2], content...", etc.

1758 Your response format should be a list of error annotations as:

1759 "[section index, error type, detailed error analysis]"

1760 If multiple errors exist in one section, list them separately.

1761 Do not output anything other than these error annotations.

1763 **Prompts K.5: LLM Prompts For Reformatting The Output**

1764 Here is the problem and the response:

1765 You are an intelligent assistant. Given a graph problem and its response, your task is to review  
1766 the response, which is divided into multiple parts with each step labeled using tags. After reading  
1767 through the steps, you should group them into distinct sections, where each section represents a  
1768 complete and logical problem-solving attempt or process.

1769 Specific instructions:

1770 1. Each section should be a standalone and complete problem-solving approach or effort.  
1771 2. For each section, include both the starting and ending tags (the ending tag should not be  
1772 earlier than the starting tag). Additionally, provide a brief summary or title of the section.  
1773 3. Present your output in this format: <<start tag>> - <<end tag>> [Brief  
1774 description]. The end tag should be the same as or later than the start tag of the  
1775 section, and the start tag of the next section should follow directly after the end tag of the  
1776 previous section.\n\n1777 Do not output any other text or explanation.

1778 Problem: {question}

1779 Response: {answer}

1780 Present your output in this format: <<start tag>> - <<end tag>> [Brief  
1781 description]. The end tag should be the same as or later than the start tag of the sec-

1782  
 1783      tion, and the start tag of the next section should follow directly after the end tag of the previous  
 1784      section. Do not output any other text or explanation.

### 1785      **Prompts K.6: LLM Prompts For Detecting The First Correct Answer**

1786  
 1787      You are a precise analysis assistant for graph theory problems. Your task is to strictly judge  
 1788      whether the provided reasoning segment contains an explicitly stated final answer that matches the  
 1789      given final answer. Given the following graph problem, a segment of reasoning process, and the  
 1790      final answer that was ultimately provided, please determine whether this segment already contains  
 1791      the final answer.

1792      The final answer is: {final\_answer}

1793      Problem: {question}

1794      Reasoning segment: {segment}

1795  
 1796      Only output YES if the final answer is explicitly presented and clearly identified in this segment  
 1797      (e.g., by phrases like "the answer is," "thus," "therefore," "in conclusion," etc.), not just mentioned  
 1798      as part of intermediate reasoning. If the final answer is not explicitly stated as a conclusion in this  
 1799      segment, output NO. Do not provide any explanation.

### 1800      **Prompts K.7: LLM Prompts For Categorizing Response Segments**

1801  
 1802      You are an expert in analyzing text relationships and functions. Your task is to strictly classify the  
 1803      function of Segment B in relation to Segment A, according to the provided definitions. Always  
 1804      output only the corresponding option without any additional explanation or commentary.

1805      Given two text segments, A and B, determine the function of B in relation to A. There are three  
 1806      possible options:

1807      1. Self-reflection: B serves to evaluate or verify the correctness of A.

1808      2. Strategy shift: B serves to alter or adjust the strategy presented in A.

1809      3. None: B has no clear functional relationship to A, or serves a different purpose not covered by  
 1810      the above categories.

1811      Output only one of the following options without any explanation: Self-reflection, Strategy shift,  
 1812      or None.

1813      Segment A: {segment\_a}

1814      Segment B: {segment\_b}

### 1815      **Prompts K.8: LLM Prompts For Identifying The Conclusive Answer Segment**

1816      Given two text segments, A and B, where B is labeled as "{label}", determine if B is an effective  
 1817      addition to A.

1818      For a segment to be effective:

- 1819      • If labeled as "Self-reflection": B should effectively evaluate or verify the correctness of A
- 1820      • If labeled as "Strategy shift": B should effectively alter or adjust the strategy presented in  
 1821      A
- 1822      • If labeled as "None": B has no clear functional relationship to A, or serves a different  
 1823      purpose not covered by the above categories.

1824      Please analyze the relationship between A and B and determine:

- 1825      1. Is B an effective addition to A given its label? (Yes/No)
- 1826      2. What is your confidence in this judgment? (High/Medium/Low)

1827      Output your judgment in the following format:

1828      Effectiveness: [Yes/No] Confidence: [High/Medium/Low]

1829      Segment A: {segment\_a}

1830      Segment B: {segment\_b}

1836  
1837

## K.2 RESPONSE REFORMATION

1838

1839

1840

1841

1842

1843

1844

1845

1846

1847

1848

1849

1850

1851

1852

1853

1854

1855

1856

1857

1858

1859

1860

1861

1862

1863

1864

1865

1866

1867

1868

1869

1870

1871

1872

1873

1874

1875

1876

1877

1878

1879

1880

1881

1882

1883

1884

1885

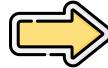
1886

1887

1888

1889

Original Response	
<p>To determine the street with the **maximum number of connections**, we need to count how many **intersections** each street is involved in.</p> <p>### Step-by-step:</p> <p>We are given a list of streets and a list of intersections (which represent connections between streets).</p> <p>#### Streets:</p> <ol style="list-style-type: none"> <li>1. Via Antonio Garbasso</li> <li>2. Via Giovanni Aldini</li> <li>3. Via Giuseppe Bagnera ... (Omitted)</li> </ol> <p>#### Intersections (connections between streets):</p> <ul style="list-style-type: none"> <li>- Via Antonio Garbasso and Via Giovanni Aldini</li> <li>- Via Antonio Garbasso and Via Giuseppe Bagnera</li> <li>- Via Antonio Garbasso and Via Orso Mario Corbino ... (Omitted)</li> </ul> <p>#### Count connections per street:</p> <ol style="list-style-type: none"> <li>1. **Via Antonio Garbasso** → **3 connections**</li> <li>2. **Via Giovanni Aldini** - Intersects with: Via Antonio Garbasso, Via Damiano Macaluso, Via Francesco Grimaldi, Via Gerolamo Cardano → **4 connections**</li> <li>3. **Via Giuseppe Bagnera** - Intersects with: Via Antonio Garbasso, Via Francesco Grimaldi, Via Gerolamo Cardano → **3 connections** ... (Omitted)</li> </ol> <p>The streets with the **maximum number of connections** are **Via Francesco Grimaldi** and **Via Gerolamo Cardano**, each with **6 connections**.</p> <p>### Final Answer:</p> <p>\boxed{6}</p>	



Reformatted Response
<p><b>[Section 1]: Introduction and Problem Setup</b> To determine the street with the **maximum number of connections**, we need to count how many **intersections** each street is involved in.</p> <p>### Step-by-step:</p> <p>We are given a list of streets and a list of intersections (which represent connections between streets).</p> <p>#### Streets:</p> <ol style="list-style-type: none"> <li>1. Via Antonio Garbasso</li> <li>2. Via Giovanni Aldini</li> <li>3. Via Giuseppe Bagnera ... (Omitted)</li> </ol> <p>#### Intersections (connections between streets):</p> <ul style="list-style-type: none"> <li>- Via Antonio Garbasso and Via Giovanni Aldini</li> <li>- Via Antonio Garbasso and Via Giuseppe Bagnera</li> <li>- Via Antonio Garbasso and Via Orso Mario Corbino ... (Omitted)</li> </ul> <p><b>[Section 2]: Counting Connections for Each Street</b> ### Count connections per street:</p> <ol style="list-style-type: none"> <li>1. **Via Antonio Garbasso** - Intersects with: Via Giovanni Aldini, Via Giuseppe Bagnera, Via Orso Mario Corbino → **3 connections**</li> <li>2. **Via Giovanni Aldini** - Intersects with: Via Antonio Garbasso, Via Damiano Macaluso, Via Francesco Grimaldi, Via Gerolamo Cardano → **4 connections** ... (Omitted)</li> </ol> <p><b>[Section 3]: Conclusion and Final Answer</b> ### Conclusion: The streets with the **maximum number of connections** are **Via Francesco Grimaldi** and **Via Gerolamo Cardano**, each with **6 connections**.</p> <p>### Final Answer:</p> <p>\boxed{6}</p>

Figure 22: Example of Response Reformation.

1890  
1891

## K.3 CASE STUDIES

1892

1893

## Case Study K.1: Response of Qwen3-32B on a DistanceK Task

1894

1895

**Problem:**

1896

1897

1898

1899

1900

1901

1902

1903

1904

1905

1906

1907

1908

1909

1910

1911

1912

1913

1914

You are given an undirected graph representing the London street network, where nodes represent streets and edges represent intersections. The distance between two directly connected nodes is 1. Given a street Westwell Road Approach, find all streets that are exactly distance 4 away from the street.

Streets in the network: Cleveland Road, Colmer Road, Danbrook Road, Streatham High Road, Penistone Road, Arragon Gardens, Ellison Road, Tankerville Road, Westwell Road, Westwell Road Approach, Buckleigh Road, Greyhound Lane, Gildersfield Road, Kempshott Road, Northanger Road Intersections in the network: Cleveland Road and Ellison Road, Colmer Road and Danbrook Road, Colmer Road and Ellison Road, Colmer Road and Streatham High Road, Danbrook Road and Ellison Road, Streatham High Road and Arragon Gardens, Streatham High Road and Greyhound Lane, Streatham High Road and Gildersfield Road, Streatham High Road and Kempshott Road, Streatham High Road and Penistone Road, Streatham High Road and Westwell Road Approach, Penistone Road and Gildersfield Road, Arragon Gardens and Ellison Road, Ellison Road and Greyhound Lane, Ellison Road and Gildersfield Road, Ellison Road and Kempshott Road, Tankerville Road and Greyhound Lane, Tankerville Road and Gildersfield Road, Tankerville Road and Kempshott Road, Westwell Road and Greyhound Lane, Westwell Road and Northanger Road, Westwell Road and Westwell Road Approach, Buckleigh Road and Greyhound Lane, Buckleigh Road and Gildersfield Road, Buckleigh Road and Kempshott Road, Buckleigh Road and Northanger Road, Greyhound Lane and Northanger Road Please find the streets in distance 4 from the street Westwell Road Approach and output the number of these streets. Present your answer in the following format: \boxed{n}. n is the number of these streets

**Response:**

## [Section 1]: Initial Understanding and Setup

The task is to find all streets exactly 4 units of distance from Westwell Road Approach in the given undirected graph, where each intersection represents an edge of distance 1. An adjacency list is a helpful way to represent the graph structure. Westwell Road Approach is the starting point (level 0). Based on the intersection list, Westwell Road Approach is directly connected to Streatham High Road and Westwell Road, making them the streets at distance 1 (level 1).

## [Section 2]: BFS Traversal to Identify Level 3 Nodes

To find the nodes at distance 2, we find the neighbors of the level 1 nodes. The neighbors of Streatham High Road are Colmer Road, Penistone Road, Gildersfield Road, Greyhound Lane, Kempshott Road, and Arragon Gardens. The neighbors of Westwell Road are Greyhound Lane and Northanger Road. The unique set of these neighbors forms level 2: Colmer Road, Penistone Road, Gildersfield Road, Greyhound Lane, Kempshott Road, Arragon Gardens, and Northanger Road. To find level 3, we find the neighbors of all level 2 nodes that have not been visited. This process yields a unique set of new nodes for level 3: Danbrook Road, Ellison Road, Tankerville Road, and Buckleigh Road.

## [Section 3]: Re-examination of Level 3 and 4 Nodes

Now, to find level 4, we must find the neighbors of the level 3 nodes (Danbrook Road, Ellison Road, Tankerville Road, Buckleigh Road) that have not already been assigned a level. The neighbors of Danbrook Road are Colmer Road (level 2) and Ellison Road (level 3), so no new nodes are found. The neighbors of Tankerville Road are Greyhound Lane, Gildersfield Road, and Kempshott Road, all of which are in level 2. The neighbors of Buckleigh Road are Greyhound Lane, Gildersfield Road, Kempshott Road, and Northanger Road, all of which are in level 2. For Ellison Road, we must check its intersections. **Ellison Road is connected to: Colmer Road and Ellison Road, Danbrook Road and Ellison Road, Arragon Gardens and Ellison Road, Ellison Road and Greyhound Lane, Ellison Road and Gildersfield Road, Ellison Road and Kempshott Road.** (This list of Ellison Road's neighbors is incomplete; it omits the intersection with Cleveland Road, which is explicitly provided in the problem description. – graph memorization error)

All of the listed neighbors for Ellison Road are in either level 2 or level 3. So no new nodes here. So indeed, no new nodes. That would mean that level 4 is empty.

(This conclusion is incorrect because the algorithm failed to process the edge between Ellison

1944 Road (level 3) and Cleveland Road. This failure to traverse a valid edge results in Cleveland Road  
 1945 being missed. – omission) [Section 4]: Final Verification and Conclusion  
 1946 The analysis incorrectly concludes that since all neighbors of level 3 nodes are already in levels 2  
 1947 or 3, there are no nodes at distance 4. The response attempts to manually trace paths but fails to  
 1948 find a path of length 4 due to the previously forgotten graph connection. **Therefore, there are no**  
 1949 **nodes at distance 4. So the answer is 0.**  
 1950 **(This final judgment is incorrect. The BFS algorithm was executed on an incomplete graph, leading**  
 1951 **to the false conclusion that level 4 is empty when it actually contains one node, Cleveland Road. –**  
 1952 **condition misjudgment)**

### Case Study K.2: Response of Qwen3-32B on a Minimum Spanning Tree Task

#### Problem:

You are required to solve the Minimum Spanning Tree Problem for an undirected street network. In this network, nodes represent streets (e.g., street IDs) and edges represent intersections between streets. The weight of each edge is the distance between two streets.

Streets in the network: Carbeen Road, Eucalyptus Drive, Erawar Close, Nicholson Avenue, Howes Close, Fairy Dell Close, Elouera Road, Booleroo Place, Higgins Place, Keys Close, Russell Crescent, Coppersleigh Close, Gundy Place, Boree Place, Brigalow Place, Baroona Road, Eastview Road, Mason Lane, Phillips Lane, Ulric Lane, Coral Heath Avenue, Brushtail Court, Hibbertia Place, Corang Road, Honeycup Close, Rocklily Avenue, Billarga Road, Pittwater Road, Apanie Place, Duneba Drive, Silver Crescent, Colin Place, Kimba Close, Western Crescent, Old Glenfield Road, Warrigal Drive, Coora Road, Namoi Road, Nulgarrar Street, Strathallen Avenue, Settlers Way, Euroka Road, Bottle Brush Road, Gum Blossom Drive, Barkala Place, Quarter Sessions Road, Dryden Avenue, De Saxe Close, Lynrob Place, Timbarra Road Intersections between these streets: Carbeen Road and Duneba Drive (weight: 3), Carbeen Road and Elouera Road (weight: 3), Eucalyptus Drive and Billarga Road (weight: 3), Eucalyptus Drive and Boree Place (weight: 9), Eucalyptus Drive and Corang Road (weight: 3), Eucalyptus Drive and Elouera Road (weight: 9), Erawar Close and Quarter Sessions Road (weight: 4), Nicholson Avenue and Quarter Sessions Road (weight: 6), Howes Close and Quarter Sessions Road (weight: 1), Fairy Dell Close and Quarter Sessions Road (weight: 5), Elouera Road and Duneba Drive (weight: 5), Booleroo Place and Corang Road (weight: 6), Higgins Place and Keys Close (weight: 3), Higgins Place and Quarter Sessions Road (weight: 4), Higgins Place and Russell Crescent (weight: 9), Russell Crescent and Quarter Sessions Road (weight: 6), Coppersleigh Close and Corang Road (weight: 2), Gundy Place and Duneba Drive (weight: 10), Boree Place and Duneba Drive (weight: 1), Brigalow Place and Duneba Drive (weight: 1), Baroona Road and Eastview Road (weight: 2), Baroona Road and Mason Lane (weight: 2), Baroona Road and Namoi Road (weight: 8), Baroona Road and Nulgarrar Street (weight: 4), Baroona Road and Phillips Lane (weight: 7), Baroona Road and Pittwater Road (weight: 9), Baroona Road and Quarter Sessions Road (weight: 9), Baroona Road and Strathallen Avenue (weight: 5), Baroona Road and Ulric Lane (weight: 1), Eastview Road and Pittwater Road (weight: 3), Eastview Road and Quarter Sessions Road (weight: 2), Mason Lane and Namoi Road (weight: 2), Coral Heath Avenue and Hibbertia Place (weight: 6), Coral Heath Avenue and Honeycup Close (weight: 7), Coral Heath Avenue and Quarter Sessions Road (weight: 1), Coral Heath Avenue and Rocklily Avenue (weight: 1), Brushtail Court and Old Glenfield Road (weight: 5), Brushtail Court and Quarter Sessions Road (weight: 3), Corang Road and Billarga Road (weight: 6), Corang Road and Kimba Close (weight: 5), Corang Road and Quarter Sessions Road (weight: 4), Pittwater Road and Quarter Sessions Road (weight: 6), Apanie Place and Duneba Drive (weight: 10), Duneba Drive and Euroka Road (weight: 6), Duneba Drive and Quarter Sessions Road (weight: 7), Silver Crescent and Coora Road (weight: 6), Silver Crescent and Quarter Sessions Road (weight: 1), Colin Place and Dryden Avenue (weight: 9), Colin Place and Quarter Sessions Road (weight: 10), Western Crescent and Quarter Sessions Road (weight: 5), Old Glenfield Road and Quarter Sessions Road (weight: 6), Warrigal Drive and Quarter Sessions Road (weight: 7), Coora Road and Barkala Place (weight: 9), Settlers Way and Quarter Sessions Road (weight: 10), Bottle Brush Road and Gum Blossom Drive (weight: 7), Bottle Brush Road and Quarter Sessions Road (weight: 4), Gum Blossom Drive and Quarter Sessions Road (weight: 10), Barkala Place and Quarter Sessions Road (weight: 7), Quarter Sessions Road and De Saxe Close (weight: 5), Quarter Sessions Road and Lynrob Place (weight: 5), Quarter Sessions Road and Timbarra Road

1998 (weight: 6), De Saxe Close and Lynrob Place (weight: 1), De Saxe Close and Timbarra Road  
 1999 (weight: 4). Identify the minimum spanning tree of this network. The minimum spanning tree is a  
 2000 subset of edges in a connected, weighted graph that connects all the vertices together with the  
 2001 smallest possible total edge weight and without any cycles. Present your answer in the following  
 2002 format:  $\boxed{n}$ , where n is the sum of the weights of the edges in the minimum spanning tree.  
 2003

**Response:**

[Section 1]: Initial Thoughts and Edge Listing

To solve this Minimum Spanning Tree (MST) problem, I will use Kruskal's algorithm, which involves sorting all edges by weight and adding them to the tree if they do not form a cycle. I will begin by listing and sorting the edges based on their weights. The edges with a weight of 1 are: Boree Place - Duneba Drive, Brigalow Place - Duneba Drive, Howes Close - Quarter Sessions Road, Baroona Road - Ulric Lane, Coral Heath Avenue - Rocklily Avenue, Coral Heath Avenue - Quarter Sessions Road, Silver Crescent - Quarter Sessions Road, and De Saxe Close - Lynrob Place. The edges with a weight of 2 are: Coppersleigh Close - Corang Road, Baroona Road - Eastview Road, Baroona Road - Mason Lane, Eastview Road - Quarter Sessions Road, and Mason Lane - Namoi Road.

[Section 2]: First Attempt at MST Calculation

I will now apply Kruskal's algorithm. First, I select all edges with weight 1, which can all be added without creating cycles. Next, I process the five edges with weight 2, all of which are safe to add. The total weight so far is 16. I then add all six edges with weight 3, bringing the total weight to 34. Next, I process edges with weight 4. **So total weight 4 edges: 4 edges? (This count is incorrect. The algorithm fails to include all edges of weight 4 from the problem description, specifically missing Higgins Place - QSR and De Saxe Close - Timbarra Road. – graph memorization error)** I add the four identified edges, bringing the total weight to 50. I then process the edges with weight 5. **So all six weight 5 edges can be added. (This count is also incorrect. The algorithm misses the edge Fairy Dell Close - QSR (weight 5). – omission)** This brings the total weight to 80.

Next, I process the edges with weight 6. **So seven edges with weight 6. (This count is incorrect. The algorithm fails to identify all edges with a weight of 6, missing several connections to Quarter Sessions Road and others. – graph memorization error)** I correctly determine that the edge Corang Road - Billarga Road should be skipped as it forms a cycle. However, I incorrectly determine that the edge Old Glenfield Road - QSR also creates a cycle based on a faulty state. **So from the seven weight 6 edges, we can add Russell Crescent - QSR, Pittwater Road - QSR, Duneba Drive - Euroka Road, Silver Crescent - Coora Road, Warrigal Drive - QSR. (This step contains multiple errors. The list of weight 6 edges is incomplete, and the cycle detection for Old Glenfield Road - QSR is based on a flawed understanding of the current graph state. – state update error)** This adds 30 to the total, making it 110. The process continues in this manner, with accumulating errors.

[Section 3]: Identifying and Correcting Mistakes in Weight 6 Edges

I now recognize that my previous analysis was flawed. I missed the edge Fairy Dell Close - QSR (weight 5). **So during the weight 5 step, I should have added seven edges instead of six. (This section demonstrates the algorithm's confusion and repeated errors. While it correctly identifies one missed edge, its subsequent recalculations remain flawed, leading to more incorrect states. This repetitive and confusing correction process is a form of redundancy. – redundancy)** I then re-evaluate the weight 6 edges and realize I also missed Booleroo Place - Corang Road, Quarter Sessions Road - Timbarra Road, and Coral Heath Avenue - Hibbertia Place. **Ah, here's another mistake! (The repeated self-correction and recalculation indicate a persistent failure to maintain a correct state of the graph's connected components. – state update error)** After several rounds of corrections, I still have an incomplete and inaccurate picture of the graph.

[Section 4]: Final MST Calculation and Verification

After realizing multiple oversights, I will restart the calculation. Weight 1 (8 edges) + Weight 2 (5 edges) + Weight 3 (6 edges) = 34. Next are the weight 4 edges. **Weight 4 edges:  $5*4=20$  – total 54. (This is incorrect. There are six edges with weight 4 in the graph. The algorithm omits the edge De Saxe Close - Timbarra Road (weight 4). – graph memorization error)** Next are the weight 5 edges. There are seven, which are all added, for a total of 35, bringing the sum to 89. Next are the weight 6 edges. **Weight 6 edges:  $9*6=54$  – total 143. (This step contains two errors. First, the algorithm incorrectly includes the edge Warrigal Drive - QSR, which has a weight of 7, not 6. – graph memorization error) (Second, this incorrect inclusion leads to a faulty sum for this step and an incorrect running total. – state update error)** Next are the weight 7 edges.

2052  
2053  
2054

Finally, to connect the remaining components, I add edges with weights 9 and 10. Total sum:  $6+10+18+20+35+54+35+9+30 = 217$ .

2055

2056

### Case Study K.3: Response of Qwen3-32B on Maximum Triangle Sum Task

2057

#### Problem:

2058

You are required to solve the Maximum Triangle Sum Problem for an undirected wikipedia network. In this network, nodes represent wikipedia articles and edges represent hyperlinks between articles. Each node is assigned a weight. Your objective is to find the triangle with the maximum sum of weights of its three nodes.

2062

Articles in the network: Montebelluna (weight: 1), Massimo Mascioletti (weight: 8), Gianluca Faliva (weight: 4), Eppelheim (weight: 7), United States (weight: 6), Alberto Rebecca (weight: 10), October 15 (weight: 6), Rugby union (weight: 1), Italy (weight: 9), List of football clubs in Italy (weight: 9), Manuel Dallan (weight: 8), Brad Johnstone (weight: 1) Hyperlinks between these articles: Montebelluna and List of football clubs in Italy, Montebelluna and Eppelheim, Montebelluna and Alberto Rebecca, Montebelluna and Italy, Montebelluna and Manuel Dallan, Massimo Mascioletti and Brad Johnstone, Massimo Mascioletti and Gianluca Faliva, Gianluca Faliva and Italy, Gianluca Faliva and Brad Johnstone, Gianluca Faliva and Manuel Dallan, Gianluca Faliva and Rugby union, Eppelheim and Italy, United States and Italy, United States and October 15, Alberto Rebecca and Italy, October 15 and Italy, October 15 and Manuel Dallan, Rugby union and Manuel Dallan, Italy and Manuel Dallan. Identify the triangle with the maximum sum of weights of its three nodes in this network. Present your answer in the following format: \boxed{n}. n is the maximum sum of weights of the triangle.

2073

2074

#### Response:

2075

##### [Section 1]: Initial Setup and Problem Understanding

2076

Okay, let's try to solve this problem. The task is to find the triangle (a set of three nodes all connected to each other) in the given undirected graph that has the maximum sum of their node weights. The nodes are Wikipedia articles with specific weights, and the edges are hyperlinks between them. First, I need to list out all the nodes and their weights to have a clear view. Let me jot them down: Montebelluna: 1, Massimo Mascioletti: 8, Gianluca Faliva: 4, Eppelheim: 7, United States: 6, Alberto Rebecca: 10, October 15: 6, Rugby union: 1, Italy: 9, List of football clubs in Italy: 9, Manuel Dallan: 8, Brad Johnstone: 1. Now, the edges (hyperlinks) are given. Let me list them out to see which nodes are connected. Each edge connects two nodes, and for a triangle, each pair in the three-node set must be connected. The edges provided are: Montebelluna - List of football clubs in Italy, Montebelluna - Eppelheim, Montebelluna - Alberto Rebecca, Montebelluna - Italy, Montebelluna - Manuel Dallan, Massimo Mascioletti - Brad Johnstone, Massimo Mascioletti - Gianluca Faliva, Gianluca Faliva - Italy, Gianluca Faliva - Brad Johnstone, Gianluca Faliva - Manuel Dallan, Gianluca Faliva - Rugby union, Eppelheim - Italy, United States - Italy, United States - October 15, Alberto Rebecca - Italy, October 15 - Italy, October 15 - Manuel Dallan, Rugby union - Manuel Dallan, Italy - Manuel Dallan.

2090

##### [Section 2]: Identifying High-Weight Nodes and Initial Triangle Candidates

2091

2092

2093

2094

2095

2096

2097

2098

2099

2100

2101

2102

2103

2104

2105

Now, I need to find all possible triangles. A triangle is three nodes where each pair is connected by an edge. So, for each set of three nodes, check if all three possible edges exist between them. But since there are 12 nodes, checking all combinations might be tedious, but maybe we can focus on nodes with higher weights because we want the maximum sum. High-weight nodes are more likely to be part of the maximum triangle. Looking at the weights, the highest are: Alberto Rebecca: 10, Italy: 9, List of football clubs in Italy: 9, Manuel Dallan: 8, Massimo Mascioletti: 8, United States: 6 (but lower than others), Others are 6, 4, 1, etc. So let's look for triangles involving Alberto Rebecca (10) and Italy (9) since those are high. Looking at the edges connected to Alberto Rebecca. From the edges, there's an edge between Montebelluna and Alberto Rebecca (edge 3), and edge 15: Alberto Rebecca - Italy. So if there's an edge between Montebelluna and Italy, then Montebelluna, Alberto Rebecca, Italy would form a triangle. Looking at edge 4: Montebelluna - Italy. Yes! So Montebelluna, Alberto Rebecca, Italy form a triangle. Their weights are  $1 + 10 + 9 = 20$ . Is there a better one? Another high-weight node is Italy (9). Let's check other nodes connected to Italy. Italy is connected to: Montebelluna (edge 4), Gianluca Faliva (edge 8), Eppelheim (edge 12), Alberto Rebecca (edge 15), United States (edge 13), Manuel Dallan (edge 19). So any triangle involving Italy and two other nodes connected to Italy. Let's see. For example, if there's a node

2106

connected to both Italy and another node connected to Italy. Let's check combinations. Take Gianluca Faliva (4) connected to Italy (edge 8). What other nodes connected to Italy are connected to Gianluca Faliva? Looking at Gianluca Faliva's edges: Edges connected to Gianluca Faliva: 7 (Massimo Mascioletti), 8 (Italy), 9 (Brad Johnstone), 10 (Manuel Dallan), 11 (Rugby union). So connected to Italy and Manuel Dallan, Brad Johnstone, etc. Now, is there a connection between Italy and Manuel Dallan? Yes, edge 19. Also, is there a connection between Gianluca Faliva and Italy (edge 8) and between Gianluca Faliva and Manuel Dallan (edge 10), and between Italy and Manuel Dallan (edge 19). So the triangle Gianluca Faliva, Italy, Manuel Dallan. Their weights:  $4 + 9 + 8 = 21$ . That's higher than 20.

2114

[Section 3]: Checking Other High-Weight Node Combinations for Better Triangles

2115

Another candidate: Italy, United States, and October 15. The edges are United States - Italy (edge 13), United States - October 15 (edge 14), and Italy - October 15 (edge 16). So that's a triangle. Their weights: 6 (US) + 9 (Italy) + 6 (October 15) = 21. Same as the previous 21. So same sum. Now, what about the triangle Gianluca Faliva, Italy, Manuel Dallan (sum 21) and the triangle United States, Italy, October 15 (sum 21). Are there any triangles with higher sums? Let's check other high-weight nodes. For example, Manuel Dallan (8) and Massimo Mascioletti (8). Are they connected? Let me check. Massimo Mascioletti is connected to Brad Johnstone (edge 6) and Gianluca Faliva (edge 7). Gianluca Faliva is connected to Manuel Dallan (edge 10). But is there a direct connection between Massimo Mascioletti and Manuel Dallan? Not that I can see. So no triangle there. What about Massimo Mascioletti (8), Gianluca Faliva (4), Brad Johnstone (1). They form a triangle? Let's check edges: Massimo Mascioletti is connected to Gianluca Faliva (edge 7) and Brad Johnstone (edge 6). Gianluca Faliva is connected to Brad Johnstone (edge 9). So yes, triangle Massimo Mascioletti, Gianluca Faliva, Brad Johnstone. Sum is  $8 + 4 + 1 = 13$ . Not better. Another high-weight node is List of football clubs in Italy (9). Let's check if it's part of a triangle. The only edge connected to it is Montebelluna (edge 1). Are there other edges? Let me check. The given edges don't show any other connections for List of football clubs in Italy. So unless Montebelluna is connected to another node that's connected to List of football clubs in Italy, but since Montebelluna is connected to List of football clubs in Italy and to Italy, but Italy isn't connected to List of football clubs in Italy? Let me check. The edges don't mention an edge between Italy and List of football clubs in Italy. So the triangle would need Montebelluna, List of football clubs in Italy, and another node connected to both. But since List of football clubs in Italy is only connected to Montebelluna, no triangle possible here. What about the triangle involving Alberto Rebecca (10), Italy (9), and someone else. For example, Alberto Rebecca is connected to Montebelluna (edge 3) and Italy (edge 15). If Montebelluna is connected to Italy (edge 4), which it is. So that triangle is already considered (sum 20). Are there others? For example, Alberto Rebecca connected to someone else who's connected to Italy. Let's see. Suppose Alberto Rebecca is connected to another node connected to Italy. But in the edges given, Alberto Rebecca is only connected to Montebelluna and Italy. So no other connections. So no other triangles involving Alberto Rebecca except with Montebelluna and Italy. Another candidate: Manuel Dallan (8), Italy (9), and someone else. Let's see. For example, if there's a node connected to both Manuel Dallan and Italy. Let's check. October 15 is connected to Italy (edge 16) and to Manuel Dallan (edge 17). So the triangle Italy, Manuel Dallan, October 15. Which we already considered (sum 21). Also, Rugby union is connected to Manuel Dallan (edge 18), but is Rugby union connected to Italy? No. So no triangle there. What about other combinations? Let's check if there's a triangle with higher than 21. Let's think about nodes with high weights. The highest weights are 10, 9, 9, 8, 8. Let's see if there's a triangle with two 9s and an 8. For example, Italy (9), List of football clubs in Italy (9), and someone connected to both. Are they connected? The edges don't mention a direct link between Italy and List of football clubs in Italy. So unless there's another node connected to both. Let's see. Montebelluna is connected to both List of football clubs in Italy (edge 1) and to Italy (edge 4). **So the triangle would be Montebelluna (1), Italy (9), List of football clubs in Italy (9).** (This is incorrect. The response correctly identifies that Montebelluna is connected to both Italy and List of football clubs in Italy, but it fails to verify the third required edge between Italy and List of football clubs in Italy before declaring it a triangle. This is a graph memorization error).

2155

Sum is  $1 + 9 + 9 = 19$ , which is less than 21.

2156

[Section 4]: Exhaustive Search for All Possible Triangles and Verification of Maximum Sum

2157

What about the triangle between Italy (9), Eppelheim (7), and someone else? For example, Eppelheim is connected to Italy (edge 12) and to Montebelluna (edge 2). But Montebelluna isn't connected to Italy except through edge 4. Wait, Montebelluna is connected to Italy (edge 4), so

2158

2159

2160 the triangle would be Montebelluna (1), Eppelheim (7), Italy (9). Sum is  $1+7+9=17$ . Not better.  
 2161 Another candidate: United States (6), Italy (9), and someone else. The only other connection  
 2162 is October 15 (edge14 and 16). So the triangle is already considered. Is there a triangle with  
 2163 three high-weight nodes? For example, Italy (9), List of football clubs in Italy (9), and someone  
 2164 else. But they aren't connected. So no. **What about Manuel Dallan (8), List of football clubs**  
 2165 **in Italy (9), and someone connected to both?** List of football clubs in Italy is only connected to  
 2166 Montebelluna (1), and Manuel Dallan is connected to Montebelluna (edge5). So the triangle would  
 2167 be Montebelluna (1), List of football clubs in Italy (9), Manuel Dallan (8). (This is incorrect. The  
 2168 response finds that Montebelluna is a common neighbor to both 'List of football clubs in Italy'  
 2169 and 'Manuel Dallan' but incorrectly assumes this forms a triangle without verifying the missing  
 2170 third edge between 'List of football clubs in Italy' and 'Manuel Dallan' – graph memorization  
 2171 error). Sum is  $1+9+8=18$ . Not better. Another possibility: Is there a triangle involving Alberto  
 2172 Rebecca (10), Manuel Dallan (8), and someone else? Let's see. Alberto Rebecca is connected  
 2173 to Montebelluna (1) and Italy (9). Manuel Dallan is connected to Italy. But there's no direct  
 2174 edge between Alberto Rebecca and Manuel Dallan. So no triangle there. What about the triangle  
 2175 between Gianluca Faliva (4), Manuel Dallan (8), and Rugby union (1)? They are connected via  
 2176 edges 10 and 18. But is there an edge between Gianluca Faliva and Rugby union? Yes (edge11).  
 2177 And between Rugby union and Manuel Dallan? Yes (edge18). And between Gianluca Faliva  
 2178 and Manuel Dallan (edge10). So the triangle is Gianluca Faliva (4), Manuel Dallan (8), Rugby  
 2179 union (1). Sum is  $4+8+1=13$ . Not better. Another check: Are there any other triangles with  
 2180 high sums? Let's think. For example, United States (6), Italy (9), and someone else. The only  
 2181 other connection is October 15 (6). So that's the triangle we already have. Sum 21. Is there a  
 2182 triangle with two 8s? For example, Massimo Mascioletti (8) and Manuel Dallan (8). But they  
 2183 aren't connected. What about Manuel Dallan (8) and someone else? Let's see. The only other  
 2184 node with 8 is Massimo Mascioletti (8). But no connection. So no triangle with two 8s. What  
 2185 about the triangle between Manuel Dallan (8), Italy (9), and another node with 8? For example,  
 2186 if there's a node connected to both Manuel Dallan and Italy. But the only such node is Gianluca  
 2187 Faliva (4) or October 15 (6). So no. So far, the maximum sum is 21, achieved by two different  
 2188 triangles: (Gianluca Faliva, Italy, Manuel Dallan) and (United States, Italy, October 15). Are  
 2189 there any other triangles with sum 21 or higher? Let me check another angle. Let's look for all  
 2190 triangles and compute their sums. First, list all possible triangles. Since this is time-consuming,  
 2191 but maybe I can find all possible triangles by looking at each node's connections. Starting with  
 2192 Italy (9), since it's connected to many nodes. Let's check all pairs of nodes connected to Italy  
 2193 and see if they are connected to each other. Italy is connected to: Montebelluna (M), Gianluca  
 2194 Faliva (G), Eppelheim (E), Alberto Rebecca (A), United States (US), October 15 (O), Manuel  
 2195 Dallan (D). So for each pair among these, check if they are connected. Check pairs: M and G:  
 2196 Are Montebelluna and Gianluca Faliva connected? No. The edges connected to Montebelluna  
 2197 are List of football clubs in Italy, Eppelheim, Alberto Rebecca, Italy, and Manuel Dallan. So no  
 2198 direct edge between M and G. M and E: Montebelluna and Eppelheim are connected (edge2). So  
 2199 the triangle M, E, Italy. Sum is  $1+7+9=17$ . M and A: Montebelluna and Alberto Rebecca are  
 2200 connected (edge3). So triangle M, A, Italy. Sum  $1+10+9=20$ . M and US: Are Montebelluna and  
 2201 United States connected? No. So no triangle. M and O: No connection between Montebelluna and  
 2202 October 15. M and D: Montebelluna and D are connected (edge5). So triangle M, D, Italy. Sum  
 2203  $1+8+9=18$ . G and E: Are Gianluca Faliva and Eppelheim connected? No. So no triangle. G and A:  
 2204 Gianluca Faliva and Alberto Rebecca? No. No edge between them. G and US: Gianluca Faliva and  
 2205 United States? No. G and O: Gianluca Faliva and October 15? No. G and D: Gianluca Faliva and  
 2206 D are connected (edge10). So triangle G, D, Italy. Sum  $4+8+9=21$ . This is one of the candidates.  
 2207 E and A: Eppelheim and Alberto Rebecca? No. No edge between them. E and US: Eppelheim  
 2208 and US? No. E and O: No. E and D: No. A and US: Alberto Rebecca and US? No. A and O: No.  
 2209 A and D: No. US and O: Yes, connected via edge14. Also, US and Italy (edge13), O and Italy  
 2210 (edge16). So triangle US, O, Italy. Sum  $6+6+9=21$ . US and D: Are US and D connected? No. So  
 2211 no triangle. O and D: Yes, connected via edge17. So triangle O, D, Italy (already considered). So  
 2212 from Italy's connections, the triangles with sum 21 are: G (4), D (8), Italy (9): sum 21, US (6),  
 2213 O (6), Italy (9): sum 21. Now check other nodes not connected to Italy. For example, check if  
 2214 there are triangles among other nodes. For example, check Gianluca Faliva's connections. He is  
 2215 connected to Massimo Mascioletti (Masc), Brad Johnstone (B), Manuel Dallan (D), Rugby union  
 2216 (R), and Italy (I). Check pairs among these: Masc and B: connected via edge6? Yes. Masc and  
 2217 Gianluca (edge7), B and Gianluca (edge9). Also, Masc and B are connected (edge6). So triangle

2214 Masc, B, Gianluca. Sum  $8+1+4=13$ . Masc and D: Are they connected? No. Masc is connected to  
 2215 Gianluca and B, but not directly to D. So no. Masc and R: No. B and D: connected via edge9 (B  
 2216 and Gianluca) and edge10 (Gianluca and D). But B and D? No. So no triangle. B and R: No. D  
 2217 and R: connected via edge18 (D and R). But Gianluca and R are connected (edge11). So triangle  
 2218 Gianluca (4), D (8), R (1). Sum  $4+8+1=13$ . Other combinations? No. Check Manuel Dallan's  
 2219 connections: Italy, October 15, Rugby union, Montebelluna. Check if any of those form triangles  
 2220 not involving Italy. For example, October 15 and Rugby union? No. October 15 and Rugby union  
 2221 are not connected. So no. Check October 15's connections: United States, Italy, Manuel Dallan.  
 2222 The triangle US-Italy-O15 is already considered. Check other nodes. For example, Massimo  
 2223 Mascioletti (8) connected to Brad Johnstone (1) and Gianluca (4). No other connections. No  
 2224 triangles. Alberto Rebecca (10) connected to Montebelluna (1) and Italy (9). No other connections.  
 2225 So no other triangles. List of football clubs in Italy is only connected to Montebelluna. So no  
 2226 triangles. Eppelheim is connected to Montebelluna and Italy. No other connections. So no other  
 2227 triangles. United States is connected to Italy and October 15. No other connections except those.  
 2228 Rugby union is connected to Gianluca and Manuel Dallan. No other connections. So all possible  
 2229 triangles have been checked. The maximum sum is 21.  
 2230  
 2231  
 2232  
 2233  
 2234  
 2235  
 2236  
 2237  
 2238  
 2239  
 2240  
 2241  
 2242  
 2243  
 2244  
 2245  
 2246  
 2247  
 2248  
 2249  
 2250  
 2251  
 2252  
 2253  
 2254  
 2255  
 2256  
 2257  
 2258  
 2259  
 2260  
 2261  
 2262  
 2263  
 2264  
 2265  
 2266  
 2267