

AlgBench: To What Extent Do Large Reasoning Models Understand Algorithms?

Henan Sun¹, Kaichi Yu³, Yuyao Wang¹, Bowen Liu¹,
Xunkai Li², Rong-Hua Li², Nuo Chen¹, Jia Li^{1,2,*},

¹The Hong Kong University of Science and Technology (Guangzhou),

²The Hong Kong University of Science and Technology,

³Beijing Institute of Technology,

magneto0617@foxmail.com, ykckevin01@gmail.com, ywang737@connect.hkust-gz.edu.cn, bliu690@connect.hkust-gz.edu.cn

cs.xunkai.li@gmail.com, lironghuabit@126.com, chenno26@gmail.com, jialee@hkust-gz.edu.cn

Abstract

Reasoning ability has become a central focus in the advancement of Large Reasoning Models (LRMs). Although notable progress has been achieved on several reasoning benchmarks such as MATH500 and LiveCodeBench, existing benchmarks for algorithmic reasoning remain limited, failing to answer a critical question: *Do LRMs truly master algorithmic reasoning?* To answer this question, we propose AlgBench, an expert-curated benchmark that evaluates LRMs under an **algorithm-centric paradigm**.

AlgBench consists of over 3,000 original problems spanning 27 algorithms, constructed by ACM algorithmic experts and organized under a comprehensive taxonomy, including Euclidean-structured, non-Euclidean-structured, non-optimized, local-optimized, global-optimized, and heuristic-optimized categories. Empirical evaluations on leading LRMs (e.g., Gemini-3-Pro, DeepSeek-v3.2-Speciale and GPT-o3) reveal substantial performance heterogeneity: while models perform well on non-optimized tasks (up to 92%), accuracy drops sharply to around 49% on globally optimized algorithms such as dynamic programming. Further analysis uncovers **strategic over-shifts**, wherein models prematurely abandon correct algorithmic designs due to **necessary low-entropy tokens**. These findings expose fundamental limitations of problem-centric reinforcement learning and highlight the necessity of an algorithm-centric training paradigm for robust algorithmic reasoning.

1 Introduction

Reasoning ability has emerged as a central concern in both the research and practical deployment of large language models (LLMs) (Chen et al., 2025, 2024b,a,c). The development of advanced reasoning-oriented models, such as

Gemini-3-pro (DeepMind, 2025), Deepseek-v3.2-speciale (Liu et al., 2025) and OpenAI’s o3 (OpenAI, 2025), underscores the significant potential of this line of inquiry. To facilitate the design of high-performance reasoning LLMs, the construction of high-quality reasoning benchmarks has become increasingly indispensable for both model training and evaluation (Glazer et al., 2024; Gulati et al., 2024; He et al., 2024; Tang et al., 2024; Huang et al., 2025, 2024).

Significant progress has been made in the construction of reasoning benchmarks, particularly within the field of mathematical reasoning, exemplified by resources such as MATH-500 (Hendrycks et al., 2021b), AIME 2024 (He et al., 2024), and AIME 2025 (Petrov et al., 2025). However, the availability of high-quality benchmarks for algorithmic reasoning remains limited, thereby hindering advancements in enhancing the algorithmic reasoning capabilities of LLMs (Li et al., 2024; Zheng et al., 2025; Yang et al., 2025; Wei et al., 2025; Niu et al., 2025; Peng et al., 2025b; Yu et al., 2025). As illustrated in Figure 1, algorithmic reasoning skills cannot be acquired merely through training on mathematical reasoning benchmarks, owing to several fundamental distinctions as follows: (1) reasoning paradigms: mathematical reasoning primarily follows a deductive paradigm, whereas algorithmic reasoning is characterized by procedural reasoning. As depicted in Figure 1, a classical example of the deductive paradigm is Aristotle’s three-stage reasoning (McLeod, 1995), which emphasizes axiom construction, theorem deduction, and precise calculation. By contrast, a canonical example of algorithmic reasoning is dynamic programming, which centers on data structuring and the systematic design of algorithmic procedures. (2) reasoning domains: mathematical reasoning is largely concerned with areas such as algebra, geometry, probability and so on, while algorithmic reasoning places greater emphasis on tasks involv-

*: Corresponding Author

ing sequence sorting, graph theory, dynamic programming and so on (Xia et al., 2025; Jain et al., 2024; Zheng et al., 2025; Peng et al., 2025a; Zhang et al., 2025a,b, 2026; Peng et al., 2026). Given these multi-level distinctions between mathematical and algorithmic reasoning, together with the current scarcity of high-quality algorithmic reasoning benchmarks, the development of such benchmarks constitutes an urgent and critical research priority.

Although several algorithmic benchmarks have been introduced, including CLRS-Text (Markeeva et al., 2024), LeetCodeDataset (Xia et al., 2025), LiveCodeBench (Jain et al., 2024) and LiveCodeBenchPro (Zheng et al., 2025), they still fail to answer a critical question: *whether LRMs truly learn algorithmic reasoning?* More specifically, these benchmarks continue to exhibit several critical limitations shown below:

Limitation 1: Existing algorithmic benchmarks frequently draw their problem sets from well-known coding platforms such as CodeForces and LeetCode. However, tasks sourced from these platforms often encompass multiple algorithms simultaneously, giving rise to several challenges. First, it becomes ill-suited for evaluating LRMs’ proficiency in a single algorithm, since failure on a multi-algorithm task does not allow researchers to isolate which specific algorithmic reasoning ability is lacking. Second, it is difficult for LLMs to acquire a well-defined algorithmic reasoning paradigm, as individual problems may require the integration of reasoning skills across several algorithms—demands that often exceed the current capabilities of LLMs.

Limitation 2: The aforementioned benchmarks encompass only a restricted subset of algorithms. For example, CLRS-Text is limited to several fundamental algorithms presented in the Introduction to Algorithms textbook (Cormen et al., 2022), while omitting a wide range of important algorithms that are commonly utilized in algorithmic competitions and industrial applications, such as tree-based dynamic programming, network flow and so on.

Limitation 3: Existing algorithmic benchmarks are highly susceptible to data contamination, as their problem sets are typically derived from publicly available sources such as textbooks and competitive programming websites. This makes it difficult to ensure that LLMs have not been exposed to these problems during pretraining (White et al.,

2024; Jain et al., 2024; Zheng et al., 2025). A parallel case has been observed in the domain of mathematical reasoning, where Wu et al. demonstrated that the Qwen2.5 series likely leveraged data from benchmarks such as MATH and AIME, thereby achieving inflated performance on mathematical reasoning tasks (Shao et al., 2025).

To address the limitations illustrated above, we introduce AlgBench, a novel benchmark for evaluating LRMs’ algorithmic reasoning. Figure 2 presents an overview of our benchmark. To summarize, our contributions are as follow:

- **Paradigm Shift:** Unlike prior benchmarks that adhere to a problem-centric paradigm, we are the first to adopt the **algorithm-centric paradigm** to the best of our knowledge, which more closely aligns with the mechanisms of human reasoning. As discussed in *Limitation 1*, existing benchmarks grounded in problem-centric paradigms are insufficient to assess whether LRMs genuinely acquire a coherent and transferable understanding of a given algorithm. To address this limitation, we adopt an algorithm-centric paradigm in the construction of our dataset, enabling a more rigorous and systematic evaluation of the algorithmic reasoning capabilities of state-of-the-art LRMs.
- **Novel-taxonomic and Contamination-free Dataset Construction:** We engage multiple algorithmic experts to manually construct problems guided by a novel algorithmic taxonomy, including Euclidean-structured, non-Euclidean-structured, non-optimized, locally optimized, globally optimized, and heuristically optimized categories. The resulting dataset comprises over 3,000 original problems spanning 27 distinct algorithms, thereby substantially enhancing algorithmic diversity while mitigating the risk of data contamination.
- **Novel Research Questions:** We systematically investigate: (1) To what extent do LRMs demonstrate genuine mastery of algorithmic reasoning? (2) What underlying factors contribute to the performance deficits of LRMs in algorithms requiring global optimization strategies? Through a comprehensive evaluation, we identify three salient findings: (1) pronounced heterogeneity in LRM performance

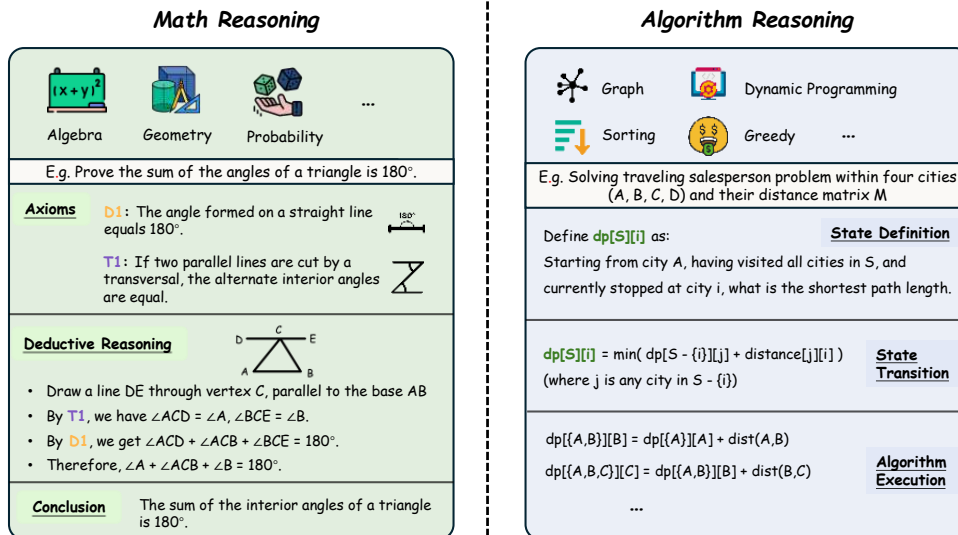


Figure 1: The difference between math reasoning and algorithmic reasoning.

across different algorithmic taxonomies; (2) uneven performance gains across these taxonomies under parameter scaling; and (3) the emergence of strategic over-shifts induced by the presence of necessary low-entropy tokens (shown in the lower-right of Figure 2).

2 Dataset Construction

Figure 2 illustrates the construction process of our dataset.

2.1 Algorithm-centric Paradigm

As discussed in Limitation 1 of Section 1, existing algorithmic benchmarks predominantly draw their problem sets from open platforms, where individual problem often requires the simultaneous application of multiple algorithmic skills. This characteristic makes it difficult to accurately assess the performance of LLMs on specific and isolated algorithms. Such benchmark pipelines follow a **problem-centric paradigm**, emphasizing the models’ ability to solve problems rather than their capacity to comprehend and internalize the underlying algorithmic principles, as illustrated in Figure 2.

In contrast, our dataset construction follows an **algorithm-centric paradigm**. Human learning does not occur solely through exposure to diverse problems; instead, individuals first acquire an understanding of algorithms and their foundational ideas, and subsequently evaluate their mastery by applying these algorithms to relevant problems. Analogously, to enable LLMs to develop genuine algorithmic reasoning abilities, we argue that

there is a pressing need for a new class of benchmarks—those designed from an algorithm-centric perspective—to evaluate model performance on individual algorithms in isolation.

It is important to note that classical algorithms are fundamentally devised to solve problems efficiently. This motivation has historically driven humans either to apply optimization techniques directly or to leverage data structures that improve the time–space complexity of computation. Consequently, the core ideas underlying classical algorithms can be broadly categorized into two groups, as illustrated in Figure 2: (1) Optimization-oriented ideas, encompassing non-optimized ideas (e.g., brute-force search), locally optimized ideas (e.g., greedy methods), globally optimized ideas (e.g., dynamic programming), and heuristic look-ahead ideas (e.g., A-star); and (2) Structure-oriented ideas, which include both Euclidean-structured and non-Euclidean-structured algorithmic principles. A detailed explanation of these categories is provided in Appendix C.1.

2.2 Dataset Curation

The detailed statistics of our constructed dataset are presented in Table 1. To address the challenges identified in *Limitation 1* and *Limitation 3* of Section 1—specifically, the issues of inaccurate evaluation of LLMs on single-targeted algorithms and potential data contamination—we engaged a group of algorithm experts. All recruited experts are accomplished ACM competitors with extensive experience and deep understanding of complex al-

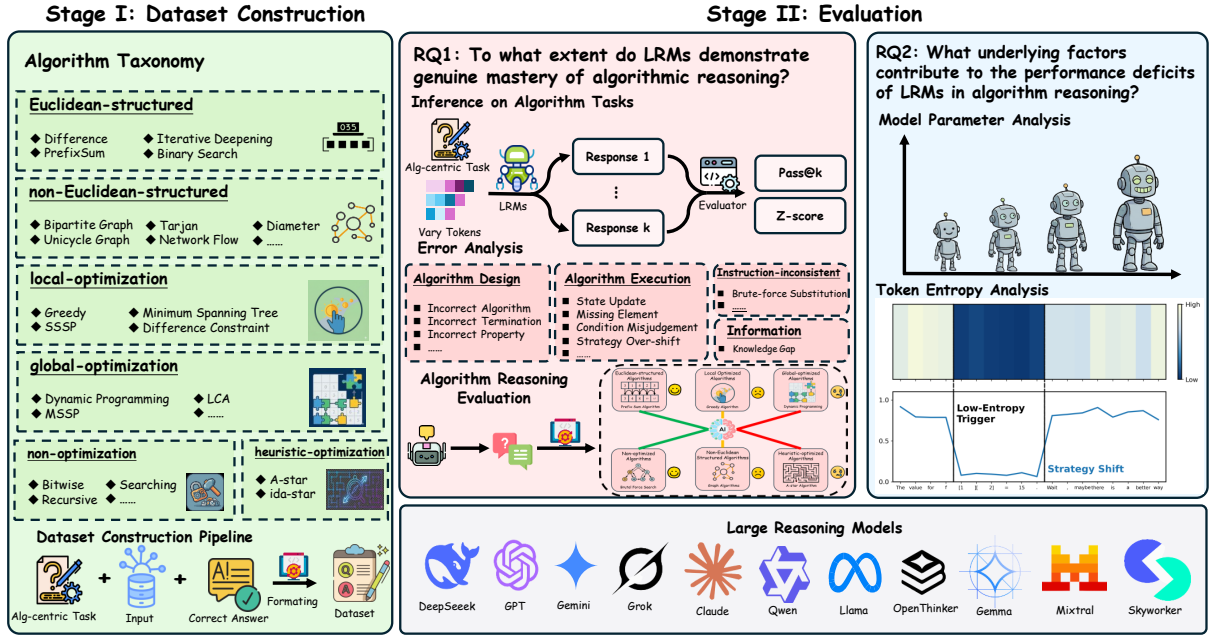


Figure 2: The overview of the proposed benchmark.

gorithmic problem solving. Under their guidance, we manually curated a collection of 3,000 original problems covering 27 distinct algorithms, each designed to require exactly one specific algorithm and verified to have no overlap with any existing open-source platforms. This construction process effectively mitigates both the evaluation bias and data contamination concerns that hinder prior benchmarks. Furthermore, in order to evaluate the mastery of LLMs on a targeted algorithm with a fine-grained perspective, we divide each problem into three difficulty levels, according to the computation complexity.

3 LLMs Evaluation

In this section, we assess the efficacy of LRMs in solving diverse algorithmic problems and critically analyze the boundaries of their reasoning capabilities. Our evaluation addresses two primary research questions: **RQ1**: To what extent do LRMs demonstrate genuine mastery of algorithmic reasoning? **RQ2**: What underlying factors contribute to the performance deficits of LRMs in algorithms requiring global optimization strategies? The experimental setup and the evaluation settings are detailed in Appendix A. Meanwhile, due to the page limit, the additional experimental results are put in Appendix D.

3.1 Exploring Robust Proficiency of LRMs in Algorithm Reasoning (RQ1)

While LRMs demonstrate generally high capabilities, their performance remains heterogeneous across distinct algorithmic domains.

Evaluating LRMs’ Ability towards Various Algorithmic Reasoning. To provide a rigorous quantitative assessment of LRM mastery, we devised a set of problems designed to isolate individual algorithm. We stratified these problems according to the taxonomy introduced in Section 2.1 and subsequently evaluated LRM efficacy across these distinct categories. For each problem instance, we sampled k responses from every comparison model and employed the $\text{Pass}@k$ metric for evaluation. Under this metric, a sample is deemed correct if at least one of the k generated responses matches the ground-truth solution. Furthermore, to mitigate the impact of disparate task difficulties across the reasoning taxonomy, we adopted the normalization protocol from prior work (Yu et al., 2023; Yuan et al., 2024). Specifically, we computed the z-score of the $\text{Pass}@k$ accuracy for each model and task, subsequently linearly rescaling these scores to the interval $[0, 1]$ as follows:

$$z_{ij} = \frac{x_{ij} - \mu(x_{i1}, \dots, x_{i|M|})}{\sigma(x_{i1}, \dots, x_{i|M|})}, \quad (1)$$

$$s_{ij} = 1.0 \cdot \frac{z_{ij} - \min(z)}{\max(z) - \min(z)},$$

Table 1: The detailed statistics of the AlgBench dataset.

Algorithm Taxonomy	Euclidean-structured				Non-Euclidean-structured							Heuristic-optimized		
Algorithms	ID	PS	Diff	BS	TDG	TUG	BGC	BGM	NF	DT	UG	AS	IDAS	
Problem Numbers	30	60	60	150	120	120	120	120	120	120	120	120	120	
Algorithm Taxonomy	Non-optimized			Local-optimized				Global-optimized						
Algorithms	BO	Recursion	Searching	Greedy	SSSP	MST	DC	LDP	IDP	TDP	BLDP	BDP	MSSP	LCA
Problem Numbers	30	150	120	120	120	120	120	150	150	120	150	150	120	120

where z_{ij} is the z-score of model i on task j , and s_{ij} is the normalized outcome.

Error Analysis. We investigate the error typologies exhibited by LRMs when solving algorithmic problems. Our methodology relies on a defined taxonomy of error categories (refer to Appendix C.2) implemented via an automated detection pipeline. Following prior work (Guo et al., 2023; Yuan et al., 2024), we first pre-process LRM outputs to enhance structural coherence. Specifically, we isolate responses containing incorrect terminal answers, segment them by double-newline delimiters, and index each resulting segment as $\langle i \rangle$. We then employ Qwen3-235B to aggregate these segments into semantic sections, summarizing each response into tuples of (start-index, end-index, summary). Finally, GPT-o3 annotates each defined section with one or more error categories, allowing us to compute the distribution of error types across the dataset.

Due to the page limit, we present the main experimental results of strong LRMs in the main body, while putting the remaining in the Appendix D. The resulting data, showed in Tables 2, Figure 3 and Figure 4, yield the following key observations:

① **LRMs cannot effectively handle complex algorithm reasoning, such as global optimization and heuristic-optimization algorithm reasoning.**

As shown in Tables 2 and Figure 3, LRMs generally exhibit superior proficiency in Euclidean-structured and non-optimization algorithms. Conversely, their performance diminishes in non-Euclidean and local-optimization contexts, with the most significant deficits observed in global-optimization and heuristic-optimized algorithms. For instance, DeepSeek-v3.2-speciale attains 88% and 92% accuracy on Euclidean-structured and non-optimization algorithms, respectively. However, performance declines moderately to 70% and 69% on non-Euclidean and local-optimization tasks, followed by a precipitous drop to 49% for both global-optimization and heuristic-optimization algorithms.

This pattern of performance degradation was consistent across the majority of evaluated models.

The primary cause of this degradation pattern largely stems from the prevailing problem-centric paradigm adopted by most reinforcement learning (RL) fine-tuning strategies for LRMs. As discussed in Section 2.1, such a paradigm inherently fails to encode the fundamental ideas underlying algorithmic reasoning, thereby limiting LRMs’ ability to acquire and execute complex algorithm reasoning. Consider the case of Dynamic Programming (DP). Under the problem-centric approach, RL fine-tuning typically aggregates datasets containing DP problem instances, corresponding solutions, and reward signals, and then directly exposes the model to these data during training. However, this process provides no mechanism for the model to infer *why* solutions to DP tasks take their specific structural form. It is well understood that a valid DP formulation must satisfy three key conditions: the presence of overlapping subproblems, the optimal substructure property, and the absence of aftereffects (i.e., Markovian structure). None of these conditions are explicitly manifested in the final solutions provided in the datasets. Consequently, LRMs are unable to reconstruct the appropriate DP states and transition functions required for systematic algorithm reasoning.

3.2 Exploring Underlying Mechanisms for Degradation Patterns of LRMs’ algorithm Reasoning (RQ2)

In addressing RQ1, we observed that LRMs exhibit heterogeneous performance across distinct algorithmic reasoning tasks. Specifically, while the majority of LRMs demonstrate proficiency in solving non-optimized and Euclidean-structured problems, they show marked limitations when confronting global- and heuristic-optimized problems. In this section, we investigate the underlying mechanisms driving this performance disparity.

Model scaling analysis. We begin by assessing

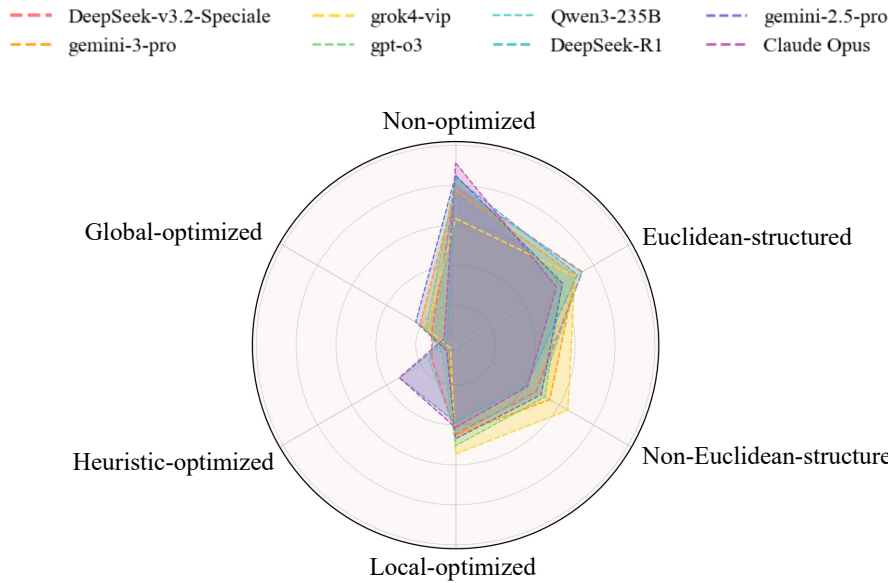


Figure 3: The normalized z-scores of LLMs on different algorithm taxonomies.

the relationship between model’s parameter and algorithmic reasoning performance. We utilize a comprehensive range of Qwen-series models for this analysis: Qwen3-4B, Qwen3-8B, Qwen3-14B, Qwen3-32B, and Qwen3-235B. Figure 5 (see in Appendix D) depicts the comparative results, detailing both the raw performance metrics and the incremental improvements observed over the base Qwen3-4B model.

Strategic over-exploration analysis. Our error analysis reveals a relatively significant prevalence of strategic over-exploration within global-optimization algorithmic reasoning tasks, as depicted in Figure 4. To elucidate the mechanisms driving this phenomenon, we calculate the policy entropy of tokens proximal to strategic shift indicators (e.g., wait, but, maybe, and so on). The details of entropy calculation and the analysis process is illustrated in Appendix C.3. The corresponding experimental results are reported from Figure 7 to Figure 11 in Appendix D.

Observing the above experimental results, we summarize the findings as follows:

② **Improvements in the algorithmic reasoning capabilities of LLMs are not commensurate with the scaling of model parameters.** As illustrated in Figure 5, the benefits of model scaling are highly non-uniform across algorithm taxonomies. LLMs exhibit substantial performance gains in local-optimized and non-Euclidean-structured rea-

soning, whereas improvements in global- and heuristic-optimized tasks are marginal. Notably, scaling yields negligible benefits for non-optimized and Euclidean-structured problems. For instance, the Qwen-series achieves relative improvements exceeding 120% and 80% in the local-optimized and non-Euclidean categories, respectively, compared to gains of less than 20% in non-optimized and Euclidean-structured domains.

Recall Finding 1 from Section 3.1, which established that LLMs exhibit superior performance in Euclidean-structured and non-optimized algorithms, but only moderate performance in local-optimized and non-Euclidean-structured contexts. Consequently, the limited scaling gains observed in the former categories can be attributed to **performance saturation**: since LLMs have already achieved high proficiency in these areas, the margin for further improvement is constrained. Conversely, non-Euclidean-structured and local-optimized algorithms are characterized by limited observation-overhead (i.e., moderate algorithmic difficulty). In these domains, the complexity is sufficient to benefit from increased capacity, which means scaling laws remain highly effective in these two algorithmic reasoning.

However, improvements in global- and heuristic-optimized algorithmic reasoning remain marginal, as well as these domains exhibiting the lowest absolute performance (see Table 2, Figure 3, and

Table 2: Leading LRMs’ performance on problems with different algorithmic ideas.

LLMs	Euclidean-structured algorithms					non-Euclidean-structured algorithms							
	ID	PS	Diff	BS	Avg.	TDG	TUG	BGC	BGM	NF	DT	UG	Avg.
DeepSeek-v3.2-Speciale	0.77	0.93	0.93	0.89	0.88	0.73	0.72	0.58	0.59	0.58	0.90	0.83	0.70
gemini-3-pro	0.93	0.93	0.48	0.95	0.83	0.86	0.58	0.56	0.62	0.63	0.90	0.78	0.70
grok4-vip	0.80	0.78	0.96	0.66	0.80	0.87	0.72	0.76	0.70	0.62	0.96	0.80	0.78
gpt-o3	0.63	0.95	0.48	0.94	0.75	0.87	0.75	0.44	0.58	0.62	0.85	0.70	0.69
Qwen3-235B	0.63	0.95	0.50	0.89	0.74	0.81	0.50	0.42	0.56	0.50	0.76	0.68	0.60
DeepSeek-R1	0.64	0.93	0.50	0.91	0.74	0.59	0.49	0.50	0.53	0.52	0.75	0.65	0.58
gemini-2.5-pro	0.63	0.88	0.50	0.55	0.64	0.67	0.53	0.37	0.58	0.48	0.80	0.48	0.56
Claude Opus	0.64	0.57	0.50	0.71	0.60	0.67	0.57	0.31	0.56	0.40	0.73	0.35	0.51
gpt-oss-120B	0.40	0.60	0.42	0.63	0.51	0.54	0.42	0.26	0.40	0.27	0.58	0.35	0.40

LLMs	local-optimized algorithms					global-optimized algorithms							
	Greedy	SSSP	MST	DC	Avg.	LDP	IDP	TDP	BLDP	BDP	MSSP	LCA	Avg.
DeepSeek-v3.2-Speciale	0.51	0.78	0.69	0.78	0.69	0.63	0.65	0.18	0.54	0.50	0.38	0.53	0.49
gemini-3-pro	0.54	0.73	0.61	0.63	0.63	0.73	0.30	0.18	0.52	0.27	0.39	0.61	0.43
grok4-vip	0.56	0.82	0.69	0.77	0.71	0.47	0.62	0.17	0.47	0.54	0.48	0.60	0.48
gpt-o3	0.52	0.78	0.64	0.78	0.68	0.59	0.47	0.20	0.56	0.47	0.40	0.48	0.45
Qwen3-235B	0.57	0.75	0.53	0.43	0.57	0.59	0.34	0.22	0.52	0.38	0.37	0.33	0.39
DeepSeek-R1	0.47	0.69	0.47	0.63	0.56	0.54	0.43	0.14	0.32	0.39	0.40	0.40	0.38
gemini-2.5-pro	0.42	0.68	0.46	0.59	0.54	0.41	0.44	0.16	0.53	0.33	0.34	0.39	0.37
Claude Opus	0.48	0.68	0.30	0.58	0.51	0.26	0.39	0.19	0.43	0.26	0.37	0.32	0.32
gpt-oss-120B	0.36	0.62	0.50	0.61	0.52	0.30	0.29	0.12	0.41	0.23	0.21	0.39	0.28

LLMs	non-optimized algorithms				heuristic-optimized algorithms		
	BO	Recursion	Searching	Avg.	AS	IDAS	Avg.
DeepSeek-v3.2-Speciale	0.97	0.93	0.87	0.92	0.39	0.58	0.49
gemini-3-pro	0.97	0.91	0.78	0.89	0.34	0.25	0.30
grok4-vip	0.90	0.53	0.87	0.77	0.40	0.39	0.39
gpt-o3	1.0	0.89	0.80	0.90	0.32	0.47	0.39
Qwen3-235B	0.80	0.78	0.77	0.78	0.38	0.45	0.41
DeepSeek-R1	0.85	0.86	0.75	0.82	0.41	0.65	0.53
gemini-2.5-pro	0.84	0.87	0.69	0.80	0.24	0.26	0.25
Claude Opus	0.85	0.89	0.61	0.79	0.37	0.55	0.46
gpt-oss-120B	0.47	0.37	0.60	0.48	0.26	0.30	0.28

Figure 5). This stagnation corroborates our observation in Section 3.1 that the prevailing problem-centric RL paradigm is insufficient for acquiring abstract algorithmic ideas. Consequently, we advocate for a paradigm shift from problem-centric to algorithm-centric RL. This transition is essential to effectively leverage scaling laws and develop LRMs with robust algorithmic reasoning capabilities.

Example 1

Let me think about this problem. We need to track both the total D and P, and the number of contestants selected, perhaps the state can be represented as $dp[i][k][d][p]$, and the transition function can be $dp[i][k][d][p] = \max \{ dp[i-1][j][m-D_j][n-P_j] \}$. Thus, the initial state is $dp[0][0][0][0] = 1$, $dp[1][0][0][0] = 0, \dots$. **Wait, maybe there is a better solution.** Let me rethink about how to model this. \dots

③ **The presence of necessary low-entropy tokens hinders the successful completion of algorithmic reasoning processes, thereby inducing multiple**

strategic over-shifts. Recall that the error analysis results presented in Figure 4 indicate a relatively high prevalence of strategic over-exploration in global-optimized algorithmic reasoning, a challenge that most LRMs struggle to address effectively. A closer examination of the false responses reveals an intriguing phenomenon: although LRMs are capable of constructing correct algorithms, they often fail to faithfully execute the designed procedures and instead prematurely abandon them in favor of alternative strategies, as shown in the Example 1.

Considering the maximum-entropy principle underlying prevailing RL paradigms, we analyze the distribution of low-entropy tokens in the vicinity of indicator tokens, as illustrated from Figure 7 to Figure 11, with a detailed illustration about token entropy calculation and analysis illustrated in Appendix C.3. The results substantiate our finding that the presence of necessary low-entropy tokens in globally optimized algorithms (e.g., numerical constants, delimiters such as “[]”, and related structural symbols) can impede the successful completion of global optimization-oriented algorithmic reasoning. This observation stands in

Qwen3-8B Error Distribution



Figure 4: Error type distributions across algorithmic reasoning taxonomies for Qwen3-8B.

tension with the dominant maximum-entropy RL paradigm, suggesting that not all low-entropy tokens should be uniformly penalized during algorithmic reasoning. Instead, RL fine-tuning could explicitly assign higher rewards to such indispensable low-entropy tokens, or alternatively, adopt agent-style paradigms that incorporate auxiliary tools to assist LRMs in executing reasoning steps when this phenomenon arises.

4 Conclusion

In this study, we presented AlgBench, a benchmark designed to move beyond problem-centric evaluation toward an algorithm-centric understanding of LRM capabilities. Through the manual curation of 3,000 original problems by algorithmic experts, we provided a contamination-free benchmark to rigorously test LRMs' mastery over specific algorithmic taxonomies.

Our comprehensive evaluation yielded three critical insights into the current state of algorithmic reasoning in LRMs:

- **Performance Heterogeneity:** LRMs exhibit superior proficiency in Euclidean-structured and non-optimized algorithms but struggle significantly with global-optimization and heuristic-optimization tasks.
- **Scaling Limitations:** Improvements in algorithmic reasoning are not commensurate with model parameter scaling. While scaling benefits algorithms like local-optimization and non-Euclidean-structured ones, it yields marginal gains for complex global- and heuristic optimization domains, suggesting that performance in these areas has reached a plateau under current training paradigms.
- **Strategic Over-shifts and Low-Entropy Barriers:** We identified that the presence of indispensable low-entropy tokens (e.g., structural delimiters and numerical constants) often triggers "strategic over-shifts". Models may construct correct algorithmic states but fail during execution, prematurely abandoning

valid strategies in favor of alternative, often incorrect, reasoning paths.

Ultimately, our results substantiate that a paradigm shift from problem-centric to algorithm-centric RL is essential for LRMs to achieve robust, human-level algorithmic reasoning.

Limitations

Although the proposed AlgBench provides an effective evaluation of the algorithmic reasoning capabilities of leading LRMs, it nevertheless exhibits several limitations, which are discussed from three perspectives below.

Complex Competitive Algorithm Gap. Although our dataset covers 27 algorithms, several complex competitive algorithms are not included in the current version (e.g., Segment Tree, Aho–Corasick and so on). Expanding the benchmark to incorporate such algorithms remains an important direction for future work.

Ambiguous Difficulty Quantification of Algorithm Taxonomies. Although our work proposes a novel algorithmic taxonomy comprising six fundamental algorithmic principles, it remains challenging to rigorously quantify the relative difficulty of these algorithms and to establish a clear relationship between algorithmic difficulty and the algorithmic reasoning capabilities of LRMs.

Limited Dataset Size. Although our dataset comprises over 3,000 problems for evaluating the algorithmic reasoning capabilities of LRMs, it remains insufficient for pre-training or fine-tuning LRMs. Consequently, continually expanding the dataset is essential to support more robust and effective pre-training or fine-tuning in future work.

Acknowledgement

This research is supported by the Guangdong Provincial Talent Program (No.2024TQ08X366).

References

- Nuo Chen, Hongguang Li, Baoyuan Wang, and Jia Li. 2024a. From good to great: Improving math reasoning with tool-augmented interleaf prompting. In *Proceedings of the 2nd Workshop on Natural Language Reasoning and Structured Explanations (@ ACL 2024)*, pages 64–79.
- Nuo Chen, Yuhan Li, Jianheng Tang, and Jia Li. 2024b. Graphwiz: An instruction-following language model for graph computational problems. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 353–364.
- Nuo Chen, Zinan Zheng, Ning Wu, Ming Gong, Dongmei Zhang, and Jia Li. 2024c. Breaking language barriers in multilingual mathematical reasoning: Insights and observations. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7001–7016.
- Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. 2025. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv preprint arXiv:2503.09567*.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2022. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. 2022. *Introduction to algorithms*. MIT press.
- Google DeepMind. 2025. A new era of intelligence with gemini 3. <https://blog.google/products/gemini/gemini-3/>. Accessed: Nov 18, 2025.
- Elliot Glazer, Ege Erdil, Tamay Besiroglu, Diego Chicharro, Evan Chen, Alex Gunning, Caroline Falkman Olsson, Jean-Stanislas Denain, Anson Ho, Emily de Oliveira Santos, and 1 others. 2024. Frontiermath: A benchmark for evaluating advanced mathematical reasoning in ai. *arXiv preprint arXiv:2411.04872*.
- Abhimanyu Gulati, Brando Miranda, Eric Chen, Emily Xia, Kai Fronsdal, Bruno de Moraes Dumont, and Sanmi Koyejo. 2024. Putnam-axiom: A functional & static benchmark for measuring higher level mathematical reasoning in llms. In *Forty-second International Conference on Machine Learning*.
- Jiayan Guo, Lun Du, Hengyu Liu, Mengyu Zhou, Xinyi He, and Shi Han. 2023. Gpt4graph: Can large language models understand graph structured data? an empirical evaluation and benchmarking. *arXiv preprint arXiv:2305.15066*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, and 1 others. 2024. Olympiadbench: A challenging benchmark for promoting agi with olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association*

- for *Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin Burns, Samir Puranik, Horace He, Dawn Song, and 1 others. 2021a. Measuring coding challenge competence with apps. [arXiv preprint arXiv:2105.09938](#).
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset. [arXiv preprint arXiv:2103.03874](#).
- Kaixuan Huang, Jiacheng Guo, Zihao Li, Xiang Ji, Jiawei Ge, Wenzhe Li, Yingqing Guo, Tianle Cai, Hui Yuan, Runzhe Wang, and 1 others. 2025. Mathperturb: Benchmarking llms’ math reasoning abilities against hard perturbations. [arXiv preprint arXiv:2502.06453](#).
- Zhen Huang, Zengzhi Wang, Shijie Xia, Xuefeng Li, Haoyang Zou, Ruijie Xu, Run-Ze Fan, Lyuman-shan Ye, Ethan Chern, Yixin Ye, and 1 others. 2024. Olympicarena: Benchmarking multi-discipline cognitive reasoning for superintelligent ai. *Advances in Neural Information Processing Systems*, 37:19209–19253.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. Live-codebench: Holistic and contamination free evaluation of large language models for code. [arXiv preprint arXiv:2403.07974](#).
- Rik Koncel-Kedziorski, Subhro Roy, Aida Amini, Nate Kushman, and Hannaneh Hajishirzi. 2016. Mawps: A math word problem repository. In *Proceedings of the 2016 conference of the north american chapter of the association for computational linguistics: human language technologies*, pages 1152–1157.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. [arXiv preprint arXiv:2302.09664](#).
- Rongao Li, Jie Fu, Bo-Wen Zhang, Tao Huang, Zhihong Sun, Chen Lyu, Guang Liu, Zhi Jin, and Ge Li. 2023. Taco: Topics in algorithmic code generation dataset. [arXiv preprint arXiv:2312.14852](#).
- Yuhan Li, Peisong Wang, Xiao Zhu, Aochuan Chen, Haiyun Jiang, Deng Cai, Victor W Chan, and Jia Li. 2024. G1bench: A comprehensive benchmark for graph with large language models. *Advances in Neural Information Processing Systems*, 37:42349–42368.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, and 1 others. 2022. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097.
- Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, and 1 others. 2025. Deepseek-v3. 2: Pushing the frontier of open large language models. [arXiv preprint arXiv:2512.02556](#).
- Larisa Markeeva, Sean McLeish, Borja Ibarz, Wilfried Bounsi, Olga Kozlova, Alex Vitvitskyi, Charles Blundell, Tom Goldstein, Avi Schwarzschild, and Petar Veličković. 2024. The clrs-text algorithmic reasoning language benchmark. [arXiv preprint arXiv:2406.04229](#).
- Owen McLeod. 1995. Aristotle’s method. *History of philosophy quarterly*, 12(1):1–18.
- Yifan Niu, Miao Peng, Nuo Chen, Yatao Bian, Tingyang Xu, and Jia Li. 2025. Reledit: Evaluating conceptual knowledge editing in language models via relational reasoning. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 10220–10238.
- OpenAI. 2025. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>. Referenced for LRM capabilities in technical benchmarking.
- Miao Peng, Nuo Chen, Zongrui Suo, and Jia Li. 2025a. Rewarding graph reasoning process makes llms more generalized reasoners. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*, pages 2257–2268.
- Miao Peng, Nuo Chen, Jianheng Tang, and Jia Li. 2025b. How does misinformation affect large language model behaviors and preferences? In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13711–13748.
- Miao Peng, Weizhou Shen, Nuo Chen, Chenliang Li, Ming Yan, and Jia Li. 2026. Incentivizing in-depth reasoning over long contexts with process advantage shaping. [arXiv preprint arXiv:2601.12465](#).
- Ivo Petrov, Jasper Dekoninck, Lyuben Baltadzhiev, Maria Drencheva, Kristian Minchev, Mislav Balunović, Nikola Jovanović, and Martin Vechev. 2025. Proof or bluff? evaluating llms on 2025 usa math olympiad. [arXiv preprint arXiv:2503.21934](#).
- Claude E Shannon. 1948. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423.
- Rulin Shao, Shuyue Stella Li, Rui Xin, Scott Geng, Yiping Wang, Sewoong Oh, Simon Shaolei Du, Nathan Lambert, Sewon Min, Ranjay Krishna, and 1 others. 2025. Spurious rewards: Rethinking training signals in rlvr. [arXiv preprint arXiv:2506.10947](#).
- Jianheng Tang, Qifan Zhang, Yuhan Li, and Jia Li. 2024. Grapharena: Benchmarking large language models on graph computational problems. [arXiv e-prints](#), pages arXiv–2407.

Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, and 1 others. 2025. Beyond the 80/20 rule: High-entropy minority tokens drive effective reinforcement learning for llm reasoning. [arXiv preprint arXiv:2506.01939](#).

Anjiang Wei, Jiannan Cao, Ran Li, Hongyu Chen, Yuhui Zhang, Ziheng Wang, Yuan Liu, Thiago SFX Teixeira, Diyi Yang, Ke Wang, and 1 others. 2025. Equibench: Benchmarking large language models’ understanding of program semantics via equivalence checking. In *NeurIPS 2025 Fourth Workshop on Deep Learning for Code*.

Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Schwartz-Ziv, Neel Jain, Khalid Saifullah, Sreemanti Dey, and 1 others. 2024. Livebench: A challenging, contamination-limited llm benchmark. [arXiv preprint arXiv:2406.19314](#).

Yunhui Xia, Wei Shen, Yan Wang, Jason Klein Liu, Huifeng Sun, Siyue Wu, Jian Hu, and Xiaolong Xu. 2025. Leetcodedataset: A temporal dataset for robust evaluation and efficient training of code llms. [arXiv preprint arXiv:2504.14655](#).

Beining Xu and Yongming Lu. 2025. Tecp: Token-entropy conformal prediction for llms. *Mathematics*, 13(20):3351.

Xin Xu, Jiabin Zhang, Tianhao Chen, Zitong Chao, Jishan Hu, and Can Yang. 2025. Ugmathbench: A diverse and dynamic benchmark for undergraduate-level mathematical reasoning with large language models. [arXiv preprint arXiv:2501.13766](#).

Lei Yang, Renren Jin, Ling Shi, Jianxiang Peng, Yue Chen, and Deyi Xiong. 2025. Probench: Benchmarking large language models in competitive programming. [arXiv preprint arXiv:2502.20868](#).

Chuanyue Yu, Kuo Zhao, Yuhan Li, Heng Chang, Mingjian Feng, Xiangzhe Jiang, Yufei Sun, Jia Li, Yuzhi Zhang, Jianxin Li, and 1 others. 2025. Graphrag-r1: Graph retrieval-augmented generation with process-constrained reinforcement learning. [arXiv preprint arXiv:2507.23581](#).

Jifan Yu, Xiaozhi Wang, Shangbin Tu, Shulin Cao, Daniel Zhang-Li, Xin Lv, Hao Peng, Zijun Yao, Xiaohan Zhang, Hanming Li, and 1 others. 2023. Kola: Carefully benchmarking world knowledge of large language models. [arXiv preprint arXiv:2306.09296](#).

Zike Yuan, Ming Liu, Hui Wang, and Bing Qin. 2024. Gracore: Benchmarking graph comprehension and complex reasoning in large language models. [arXiv preprint arXiv:2407.02936](#).

Qifan Zhang, Nuo Chen, Zehua Li, Miao Peng, Jing Tang, and Jia Li. 2025a. Improving llms’ generalized reasoning abilities by graph problems. [arXiv preprint arXiv:2507.17168](#).

Qifan Zhang, Xiaobin Hong, Jianheng Tang, Nuo Chen, Yuhan Li, Wenzhong Li, Jing Tang, and Jia Li. 2025b. Gcoder: Improving large language model for generalized graph reasoning. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 4149–4159.

Qifan Zhang, Jianhao Ruan, Aochuan Chen, Kang Zeng, Nuo Chen, Jing Tang, and Jia Li. 2026. Exposing weaknesses of large reasoning models through graph algorithm problems. [arXiv preprint arXiv:2602.06319](#).

Zihan Zheng, Zerui Cheng, Zeyu Shen, Shang Zhou, Kaiyuan Liu, Hansen He, Dongruixuan Li, Stanley Wei, Hangyi Hao, Jianzhu Yao, and 1 others. 2025. Livecodebench pro: How do olympiad medalists judge llms in competitive programming? [arXiv preprint arXiv:2506.11928](#).

A Experimental Setup

Evaluated LRMs. Our evaluation encompasses 25 state-of-the-art LRMs, including both closed-source commercial models and open-source LRMs. Leveraging AlgBench, we conduct a comprehensive assessment of the algorithmic reasoning capabilities of contemporary LRMs. The evaluated models are listed in Table 3:

Table 3: The detailed information about evaluated LRMs.

Model Name	Organization	Release Date
DeepSeek-v3.2-Special	DeepSeek	December 2025
gemini-3-pro	Google DeepMind	November 2025
grok4	xAI	July 2025
gpt-o3	OpenAI	April 2025
Qwen3-235B	Alibaba	July 2025
DeepSeek-R1	DeepSeek	January 2025
gemini-2.5-pro	Google DeepMind	April 2025
Claude Opus	Anthropic	November 2025
gpt-oss-120B	OpenAI	August 2025
Qwen3-4B	Alibaba	April 2025
Qwen3-8B	Alibaba	April 2025
Qwen3-14B	Alibaba	April 2025
Qwen3-32B	Alibaba	April 2025
QwQ-32B	Alibaba	March 2025
DeepSeek-R1-0528-Qwen3-8B	DeepSeek	May 2025
DeepSeek-R1-Distill-Qwen-32B	DeepSeek	May 2025
OpenThinker2-32B	OpenThoughts	March 2025
Skyworker-OR1-32B	SkyworkAI	May 2025
Qwen3-Coder-30B-A3B-Instruct	Alibaba	July 2025
Llama-3.3-70b-instruct	Meta	December 2024
gemma3-12B	Google	March 2025
gemma3-27B	Google	March 2025
Llama3-8B	Meta	September 2024
Llama3.1-8B	Meta	September 2024
Mixtral-8x7B	Mistral AI	December 2023

Evaluation Settings. We employ $\text{pass}@k$ to evaluate the average performance of LRMs on algorithmic reasoning, and the z-score to mitigate the impact of disparate problems’ difficulties across the algorithmic reasoning taxonomy, as illustrated in Section 3.1. To evaluate the performance of LRMs

on a target algorithm, each prompt is appended with explicit instructions (e.g., “Please solve this problem using the XXX algorithm”). The detailed prompt templates are provided in Appendix C.4.

To maintain consistency in evaluations and facilitate reproduction, we set the maximum output length to 32k tokens and employ a greedy decoding strategy with temperature 0.6.

Experiment Environment. To ensure reproducibility, we report the hardware and software configurations employed in our experiments. All experiments were conducted on a server equipped with an Intel(R) Xeon(R) Gold 6240 CPU (2.60 GHz) and 8 NVIDIA A800 GPUs, each with 80 GB of memory, using CUDA version 12.4.

B Related Work

Benchmarking LRMs on Math Problems. Mathematical reasoning has long attracted many attentions from both academic and industry. Multiple math-related benchmarks have been proposed in this area (Koncel-Kedziorski et al., 2016; Hendrycks et al., 2021b; Cobbe et al., 2021; Chen et al., 2022; He et al., 2024; Xu et al., 2025; Petrov et al., 2025). Among these, MATH-500 (Hendrycks et al., 2021b), AIME 2024 (He et al., 2024), and AIME 2025 (Petrov et al., 2025) are those among the most representative datasets for the training and evaluation of LRMs on high-level mathematical reasoning abilities. However, as stated in Section 1, these training and evaluation cannot help improve or testify the reasoning capacity of LRMs on algorithm reasoning, since there are significant difference between these two reasoning domains. Thus, in order to thoroughly promote the development of LRMs towards genuine artificial general intelligence (AGI), it is necessary for building benchmarks for LRMs’ algorithmic reasoning abilities as well.

Benchmarking LRMs on Competitive Algorithmic Programming. As algorithmic reasoning constitutes a fundamental capability that LRMs must acquire to advance toward genuine AGI, a number of benchmarks have been proposed to evaluate such abilities (Hendrycks et al., 2021a; Li et al., 2022, 2023; Xia et al., 2025; White et al., 2024; Jain et al., 2024; Zheng et al., 2025). Among these, several benchmarks achieve considerable influence among academic and industrial areas, such as LeetCode-Dataset (Xia et al., 2025), LiveCodeBench (Jain et al., 2024) and LiveCodeBench Pro (Zheng et al.,

2025). However, all of these existing benchmarks predominantly adopt a problem-centric paradigm, emphasizing the accuracy of LRMs in solving algorithmic tasks rather than assessing whether they have genuinely internalized the underlying principles of a specific algorithm. To achieve this objective, and to the best of our knowledge, we are the first to introduce a paradigm shift toward an algorithm-centric framework, which places primary emphasis on assessing whether LRMs have genuinely mastered an algorithm and its underlying core principles.

C Definition and Explanation

C.1 Detailed Illustration of Algorithm Taxonomy

This appendix provides a detailed rationale for the taxonomy used in the dataset. The algorithms are categorized based on the **algorithm-centric paradigm**.

C.1.1 Euclidean-structured

Category Characteristics: This category encompasses algorithms that operate on data structures with a defined linear order or coordinate system (e.g., arrays, sequences, grids). The core feature is that the data elements have fixed, indexable positions. Problem-solving relies heavily on dimension partitioning, interval operations, or searching within a linear metric space.

Algorithm Justification:

Difference Arrays & Prefix Sum: Both algorithms strictly depend on the continuity and ordering of linear arrays. They utilize the sequential arrangement of elements to perform range updates or queries in $O(1)$ time, leveraging the Euclidean property of distance and position.

Binary Search: This algorithm is predicated on the "sorted" property, which implies a strong linear constraint. It operates in a 1D Euclidean space (a number line), systematically reducing the search interval based on coordinate comparison.

Iterative Deepening: While often applied to graphs, the mechanism itself maps a complex search space onto a linear dimension (depth $d = 1, 2, 3 \dots$). It treats the search radius as a gradually expanding Euclidean boundary.

C.1.2 Non-Euclidean-structured

Category Characteristics: Algorithms in this category deal with topological structures where elements (nodes) are connected by relations (edges) rather than fixed coordinates. The focus is on connectivity, structural shape (cycles, components), and flow, rather than spatial position.

Algorithm Justification:

Bipartite Graph, Diameter, & Unicyclic: These tasks involve analyzing the intrinsic shape of the graph. Identifying a bipartite set, finding the longest path (diameter), or detecting a ring in a tree (unicyclic) are purely topological properties independent of embedding.

Connectivity (Directed/Undirected):

Connectivity algorithms (finding SCCs or BCCs) analyze the reachability among nodes, a fundamental property of non-Euclidean graph theory.

Network Flow & Bipartite Matching: These are complex graph problems that view edges as capacities or relationships. They optimize flow or associations based strictly on the graph's adjacency matrix, not on geometric proximity.

C.1.3 Non-optimized

Category Characteristics: These algorithms are fundamentally procedural or computational in nature. Rather than incorporating optimization principles to reduce time or space complexity, they rely on deterministic computations, logical operations, or exhaustive traversal strategies to produce correct outputs.

Algorithm Justification:

Bitwise: Involves low-level arithmetic and logical manipulation of binary representations. It is a calculation process, not an optimization search.

Recursive: Represents a control flow mechanism for task decomposition. Without an associated cost function or memoization, pure recursion is simply a method of execution.

Searching (BFS/DFS): Fundamental traversal strategies. Used alone, the goal is visiting nodes or verifying existence, not optimizing a numerical value.

C.1.4 Local-optimized

Category Characteristics: This category encompasses algorithms that iteratively construct solutions by selecting locally optimal actions at each step, such as those satisfying the greedy-choice property, or by applying local relaxation operations. Such algorithms generally do not preserve a comprehensive global state history and are predicated on the assumption that a sequence of local improvements will converge to a globally optimal solution.

Algorithm Justification:

Greedy: Greedy algorithm is one of the most classical local-optimized algorithms. It makes irrevocable decisions only based on locally optimal strategies without observing the global state information.

Shortest Single (Dijkstra/Bellman-Ford):

These algorithms are based on relaxation, a local operation that evaluates whether a candidate path can be improved through a given edge. In particular, Dijkstra's algorithm adopts a greedy strategy by iteratively selecting the unvisited node with the minimum tentative distance.

MST (Prim/Kruskal): The construction of a Minimum Spanning Tree is a canonical example of a greedy paradigm. Algorithms such as Prim's and Kruskal's incrementally expand the spanning structure by selecting, respectively, the nearest admissible node or the minimum-weight edge. These selection rules are inherently locally optimal, as they rely exclusively on immediate, local information rather than on a globally optimized view of the graph.

Difference Constraints: These systems are typically solved by transforming them into Single-Source Shortest Path problems, thus inheriting the local relaxation characteristics of Bellman-Ford or SPFA.

C.1.5 Global-optimization

Category Characteristics: Algorithms in this category address problems by decomposing them into overlapping subproblems and systematically integrating their solutions. A defining characteristic is the maintenance of a global state—typically

through dynamic programming tables or memoization—which guarantees that the final solution is globally optimal by accounting for all relevant state transitions.

Algorithm Justification:

DP Variants (Linear, Interval, Tree, Bitmask and Binary Lifting): Dynamic Programming is the epitome of global optimization. It systematically explores all relevant states and transitions, using the valid DP transition function to guarantee the global optimum is found from sub-solutions.

LCA: Although primarily employed to answer queries, the preprocessing stage constructs a global state table (e.g., $ancestor[i][j]$, denoting the 2^j -th ancestor of node i). This procedure is governed by the state transition relation $f(i, j) = f(f(i, j - 1), j - 1)$, which exemplifies the global state propagation and optimization logic characteristic of dynamic programming.

Shortest Multi (Floyd-Warshall): In contrast to Dijkstra’s algorithm, Floyd–Warshall adopts a dynamic programming framework that systematically accounts for the effect of every node as a potential intermediate vertex between all pairs of nodes, thereby achieving a globally optimized solution over the entire distance matrix.

C.1.6 Heuristic-optimization

Category Characteristics: This category serves as an intermediate paradigm between deterministic optimization methods and heuristic search strategies. It incorporates domain-specific knowledge through a heuristic evaluation function $f(n) = g(n) + h(n)$, which combines the accumulated cost of the current state with an estimate of the remaining cost to the goal. By jointly accounting for present progress and future potential, this formulation guides and prioritizes the search process, enabling more efficient solution discovery.

Algorithm Justification:

A* (A-Star): One of the most classical algorithms for heuristic search. It optimizes the path-finding process by adding an estimated future cost $h(n)$ to the known path cost, guiding the search intelligently.

IDA*: Combines the space efficiency of Iterative Deepening with the directional guidance of heuristics. It uses the heuristic threshold to prune branches that exceed the estimated optimal cost.

C.2 Error Type Classification

We categorize model errors into four major classes:

C.2.1 Algorithm Design (AD)

Errors in this category arise when the model fails to formulate a theoretically sound algorithmic approach for the given problem.

AV (Admissibility Violation) The heuristic function used in search algorithms (e.g., A*) overestimates the true cost to the goal ($h(n) > actual_cost$), invalidating optimality guarantees or causing the pruning of optimal paths.

CCV (Capacity Constraint Violation) In network flow problems, the model proposes a flow assignment that exceeds edge capacities or results in invalid negative flows.

DMC (Direct Mathematical Calculation) The model attempts to bypass the required algorithmic design (e.g., Dynamic Programming) by deriving the answer directly through a mathematical formula or shortcut.

EFE (Evaluation Function Error) The total estimated cost $f(n)$ is calculated incorrectly (e.g., $f(n) \neq g(n) + h(n)$) within heuristic search algorithms.

FCV (Flow Conservation Violation) The model violates flow conservation constraints (inflow \neq outflow) for non-source/sink nodes in network flow problems.

IBMS (Initial Bounds Mis-specified) In binary search, the starting bounds (low/high) fail to cover the feasible range or exclude the correct answer.

IA (Incorrect Algorithm) The model selects an algorithm that is fundamentally unsuitable for the problem constraints or objectives (e.g., using a greedy approach for a problem requiring DP).

IGS (Incorrect Greedy Strategy) The model identifies a flawed local criterion for making greedy choices, leading to suboptimal or

invalid solutions, often failing the exchange argument.

IH (Ineffective Heuristic) The heuristic is technically admissible (e.g., $h(n) = 0$) but provides poor guidance, causing the search to degrade into a less efficient algorithm (e.g., A* becoming Dijkstra's)12.

MCLE (Midpoint Comparison Logic Error)

The comparison logic in binary search is inverted or incorrect, leading to search directions that exclude the target.

MPV (Markov Property Violation) In Dynamic Programming, the state definition fails to encapsulate all necessary history, leading to invalid transitions.

OSV (Optimal Substructure Violation) The model fails to recognize that the optimal solution can be constructed from optimal sub-solutions, leading to incorrect DP state definitions or recurrence relations.

PM (Problem Misunderstanding) The model misinterprets the problem's constraints, objective function, or edge cases, resulting in a flawed algorithmic model.

SPAЕ (Shortest Path Algorithm Error) In difference constraints or graph problems, the model selects an inappropriate shortest path algorithm (e.g., using an algorithm that cannot handle negative weights when they are present).

SSPV (Search Space Precondition Violation) The model applies an algorithm (e.g., binary search) without verifying essential preconditions, such as array sortedness or monotonicity.

SSRE (State Space Representation Error) The definition of the search space is flawed, including incomplete state definitions, incorrect successor generation, or miscalculated edge costs.

C.2.2 Algorithm Execution (AE)

Errors in this category occur when the high-level logic is generally correct, but mistakes are made during the step-by-step implementation or calculation of the algorithm.

AIE (Ancestor Identification Error) In Lowest Common Ancestor (LCA) problems, the model fails to correctly identify the common ancestor or finds one that is not the lowest.

APE (Augmenting Path Error) The model fails to correctly identify or construct augmenting paths in matching or flow problems, or terminates the search prematurely.

BIE (Bridge Identification Error) In graph connectivity, the model uses incorrect conditions (e.g., comparing $low[child]$ and $dfn[current]$) to identify bridge edges.

BMC (Bipartite Matching Confusion) The model misapplies theorems (e.g., König's theorem) regarding the relationship between vertex covers, independent sets, and matchings.

BPE (Bipartite Partition Error) The model incorrectly assigns vertices to partitions or fails to identify the bipartite structure of the graph.

BVE (Bipartite Verification Error) The model incorrectly determines graph bipartiteness or assumes it without verification.

CDE (Cycle Detection Error) The model fails to detect cycles during graph construction (e.g., in MST algorithms), allowing invalid edges to be added.

CE (Computation Errors) General arithmetic mistakes occurring during intermediate calculations, prefix sum accumulations, or variable updates.

CNIE (Cycle Node Identification Error) In unicyclic graphs, the model confuses nodes belonging to the cycle with those in attached tree branches.

CRE (Complement Relationship Error) The model fails to correctly apply the complement relationship (e.g., $|MaxIndependentSet| = |V| - |MinVertexCover|$).

CVIE (Cut Vertex Identification Error) The model incorrectly identifies articulation points, often mishandling root node cases or low-link value conditions.

ESE (Edge Sorting Error) In algorithms like Kruskal's, the model fails to sort edges correctly by weight.

FACE (Final Answer Construction Error) The algorithmic execution is sound, but the model fails to format, index, or construct the final output correctly.

GCE (Graph Construction Error) The model makes mistakes in building the underlying graph structure, such as missing edges or incorrect node representations.

IC (Index Confusions) The model confuses 0-based and 1-based indexing, or mixes row/column indices, leading to off-by-one errors.

IT (Incorrect Termination) The algorithm terminates at the wrong moment (e.g., A* stopping when a goal is seen rather than popped from the priority queue).

NPE (Neighbor Processing Error) The model incorrectly identifies valid neighbors, fails boundary checks, or processes neighbors in an invalid order during graph traversal.

PQE (Priority Queue Error) The model mismanages priority queue operations, such as incorrect priority updates or extraction logic.

PTE (Path Traversal Error) The model fails to correctly navigate paths in trees or graphs, missing nodes or making wrong directional decisions.

ROE (Relaxation Operation Error) In shortest path algorithms, the model fails to correctly implement the edge relaxation step or update distances.

SCCEE (SCC Extraction Error) In Tarjan's algorithm, the model fails to correctly identify the condition for extracting an SCC from the stack, where SCC stands for Strongly Connected Component.

SOE (Strategic Over-exploration) In Dynamic Programming, the model establishes correct states but wastes computation searching for unnecessary "better" strategies instead of executing the transition.

STE (State Transition Error) The model applies the DP transition formula incorrectly, using wrong indices or logic during the update step.

TBPE (Tree Branch Processing Error) In unicyclic graphs, the model mishandles the processing of tree components attached to cycle nodes.

TTE (Tree Traversal Error) The model implements incorrect logic for DFS/BFS traversals on trees, such as visiting nodes redundantly.

C.2.3 Instruction-inconsistent Algorithm (IIA)

Errors where the model deviates from the specific algorithmic method requested in the prompt, regardless of the correctness of the alternative solution.

BFS (Brute-force Substitution) The model resorts to exhaustive enumeration or nested loops instead of the requested efficient algorithm (e.g., replacing DP or binary search with linear scans).

DMC (Direct Mathematical Calculation) The model skips the requested algorithmic process (e.g., DP or simulation) entirely in favor of a direct formula derivation.

EU (Execution Unwillingness) The model explicitly refuses to perform the detailed steps of the algorithm (e.g., refusing to compute a prefix sum matrix).

RG (Random Guess) The model abandons logical reasoning and provides a baseless answer or switches to an unrelated method without justification.

C.2.4 Information (Info)

KG (Knowledge Gap) The model demonstrates a fundamental lack of understanding of key concepts (e.g., optimal substructure, Markov property, or specific graph theorems) required to solve the problem.

C.3 Detailed Illustration of Token Entropy Analysis

Token Entropy Calculation. We follow the prior work of entropy calculation and obtain the corresponding token entropy in our experiments (Shannon, 1948; Kuhn et al., 2023; Wang et al., 2025; Xu and Lu, 2025). Formally, given the policy distribution p_i at token index i and the vocabulary V , the corresponding token entropy $H(p_i)$ is defined as follow:

$$H(p_i) = - \sum_{j=0}^{|V|} p_i(v_j) \log p_i(v_j). \quad (2)$$

Token Entropy Analysis. After computing token entropy, we rank all tokens in each LRM-generated response according to their entropy values and identify the corresponding top- k high-entropy tokens and bottom- k low-entropy tokens. We then quantify the co-occurrence between indicator tokens (e.g., wait, but, and so on) and the selected high-/low-entropy tokens by counting their occurrences within a predefined local context window around each indicator token (specifically, a distance of 10 tokens in our experiments). Based on this procedure, we obtain the experimental results presented in Figures 7 through 11.

C.4 Additional Algorithm Question Explanation

This section illustrates several examples of our prompts about the algorithm questions, including the system prompts and the question prompts, as shown in the Example 2 and 3 below.

Example 2: System Prompt

You are a professional algorithm assistant. Please answer questions according to the following requirements:

1. Do not use any code execution or programming tools to calculate the answer.
2. You must give your thinking process before outputting the final result in the format "Final answer: [number1, number2, number3, ...]".
3. Make sure the output array format is correct, with numbers separated by commas and spaces.

Note: Do not rely on code execution results.

Example 3: Question Prompt

There are 8 students ready to take photos, arranged in 3 rows with [6, 1, 1] persons respectively. The first row is in the front, and the last row is in the back. The students have distinct heights and are labeled from 1 to 8 in descending order of height. During the photo arrangement, it is required that in each row, the heights decrease from left

to right, and in each column, the heights decrease from back to front. How many different valid arrangements are possible for the group photo? **Please solve this using a linear dynamic programming algorithm.**

As illustrated in Section 1 and Appendix A, one of the primary goals of this benchmark is to evaluate LRMs' algorithmic reasoning ability towards a targeted algorithm. Thus, we specifically require LRMs not to use any code execution in system prompt and also require LRMs to use the specified algorithm in the question prompt.

D Additional Experimental Results

This section presents the additional experimental results in Section 3, including Table 4, Table 5, Figure 5, Figure 6, Figure 7, Figure 8, Figure 9, Figure 10 and Figure 11.

Table 4: Complete Models’ performance on algorithms of structure-oriented, local-optimized and global-optimized ideas.

LLMs	Euclidean-structured algorithms					non-Euclidean-structured algorithms							
	ID	PS	Diff	BS	Avg.	TDG	TUG	BGC	BGM	NF	DT	UG	Avg.
DeepSeek-v3.2-Speciale	0.77	0.93	0.93	0.89	0.88	0.73	0.72	0.58	0.59	0.58	0.90	0.83	0.70
gemini-3-pro	0.93	0.93	0.48	0.95	0.83	0.86	0.58	0.56	0.62	0.63	0.90	0.78	0.70
grok4-vip	0.80	0.78	0.96	0.66	0.80	0.87	0.72	0.76	0.70	0.62	0.96	0.80	0.78
gpt-o3	0.63	0.95	0.48	0.94	0.75	0.87	0.75	0.44	0.58	0.62	0.85	0.70	0.69
Qwen3-235B	0.63	0.95	0.50	0.89	0.74	0.81	0.50	0.42	0.56	0.50	0.76	0.68	0.60
DeepSeek-R1	0.64	0.93	0.50	0.91	0.74	0.59	0.49	0.50	0.53	0.52	0.75	0.65	0.58
gemini-2.5-pro	0.63	0.88	0.50	0.55	0.64	0.67	0.53	0.37	0.58	0.48	0.80	0.48	0.56
Claude Opus	0.64	0.57	0.50	0.71	0.60	0.67	0.57	0.31	0.56	0.40	0.73	0.35	0.51
gpt-oss-120B	0.40	0.60	0.42	0.63	0.51	0.54	0.42	0.26	0.40	0.27	0.58	0.35	0.40
Qwen3-4B	0.60	0.93	0.40	0.62	0.64	0.24	0.29	0.23	0.38	0.26	0.32	0.13	0.26
Qwen3-8B	0.53	0.93	0.47	0.76	0.67	0.33	0.43	0.32	0.46	0.28	0.39	0.23	0.35
Qwen3-14B	0.57	0.93	0.50	0.83	0.71	0.45	0.43	0.38	0.48	0.33	0.47	0.27	0.40
Qwen3-32B	0.67	0.45	0.38	0.83	0.58	0.50	0.48	0.29	0.48	0.33	0.38	0.27	0.39
QwQ-32B	0.57	0.95	0.37	0.83	0.68	0.62	0.40	0.33	0.50	0.41	0.62	0.34	0.46
DeepSeek-R1-0528-Qwen3-8B	0.50	0.88	0.40	0.73	0.63	0.38	0.45	0.28	0.43	0.34	0.46	0.28	0.37
Distill-Qwen-32B	0.50	0.95	0.45	0.76	0.67	0.50	0.39	0.19	0.47	0.38	0.45	0.22	0.37
OpenThinker2-32B	0.63	0.80	0.50	0.81	0.69	0.43	0.41	0.28	0.43	0.40	0.55	0.26	0.39
Skyworker-OR1-32B	0.67	0.93	0.50	0.86	0.74	0.56	0.49	0.33	0.49	0.40	0.58	0.29	0.45
Qwen3-Coder-30B-A3B-Instruct	0.47	0.80	0.22	0.57	0.51	0.12	0.12	0.08	0.39	0.17	0.21	0.07	0.17
Llama-3.3-70b-instruct	0.17	0.62	0.12	0.38	0.32	0.07	0.08	0.03	0.17	0.03	0.19	0.10	0.09
gemma3-12B	0.33	0.40	0.0	0.39	0.28	0.10	0.03	0.09	0.30	0.04	0.18	0.04	0.11
gemma3-27B	0.30	0.50	0.08	0.49	0.34	0.11	0.08	0.06	0.29	0.08	0.22	0.08	0.13
Llama3-8B	0.07	0.07	0.0	0.13	0.07	0.05	0.03	0.07	0.02	0.03	0.09	0.03	0.04
Llama3.1-8B	0.07	0.13	0.0	0.20	0.10	0.03	0.03	0.06	0.09	0.02	0.08	0.03	0.05
Mixtral-8x7B	0.07	0.03	0.0	0.15	0.06	0.08	0.01	0.01	0.09	0.0	0.13	0.02	0.05
LLMs	local-optimized algorithms					global-optimized algorithms							
	Greedy	SSSP	MST	DC	Avg.	LDP	IDP	TDP	BLDP	BDP	MSSP	LCA	Avg.
DeepSeek-v3.2-Speciale	0.51	0.78	0.69	0.78	0.69	0.63	0.65	0.18	0.54	0.50	0.38	0.53	0.49
gemini-3-pro	0.54	0.73	0.61	0.63	0.63	0.73	0.30	0.18	0.52	0.27	0.39	0.61	0.43
grok4-vip	0.56	0.82	0.69	0.77	0.71	0.47	0.62	0.17	0.47	0.54	0.48	0.60	0.48
gpt-o3	0.52	0.78	0.64	0.78	0.68	0.59	0.47	0.20	0.56	0.47	0.40	0.48	0.45
Qwen3-235B	0.57	0.75	0.53	0.43	0.57	0.59	0.34	0.22	0.52	0.38	0.37	0.33	0.39
DeepSeek-R1	0.47	0.69	0.47	0.63	0.56	0.54	0.43	0.14	0.32	0.39	0.40	0.40	0.38
gemini-2.5-pro	0.42	0.68	0.46	0.59	0.54	0.41	0.44	0.16	0.53	0.33	0.34	0.39	0.37
Claude Opus	0.48	0.68	0.30	0.58	0.51	0.26	0.39	0.19	0.43	0.26	0.37	0.32	0.32
gpt-oss-120B	0.36	0.62	0.50	0.61	0.52	0.30	0.29	0.12	0.41	0.23	0.21	0.39	0.28
Qwen3-4B	0.27	0.49	0.17	0.33	0.31	0.22	0.27	0.22	0.31	0.18	0.28	0.27	0.25
Qwen3-8B	0.32	0.51	0.26	0.38	0.36	0.37	0.27	0.30	0.32	0.22	0.28	0.31	0.29
Qwen3-14B	0.39	0.45	0.29	0.38	0.38	0.45	0.31	0.28	0.36	0.28	0.37	0.34	0.34
Qwen3-32B	0.34	0.52	0.22	0.18	0.32	0.35	0.35	0.27	0.33	0.17	0.38	0.40	0.32
QwQ-32B	0.33	0.58	0.19	0.53	0.41	0.49	0.31	0.19	0.40	0.28	0.33	0.37	0.34
DeepSeek-R1-0528-Qwen3-8B	0.25	0.52	0.16	0.53	0.36	0.34	0.21	0.19	0.34	0.23	0.28	0.28	0.27
Distill-Qwen-32B	0.27	0.43	0.24	0.42	0.34	0.37	0.27	0.09	0.31	0.20	0.29	0.31	0.26
OpenThinker2-32B	0.31	0.57	0.25	0.40	0.38	0.33	0.22	0.11	0.35	0.23	0.32	0.36	0.27
Skyworker-OR1-32B	0.33	0.50	0.38	0.48	0.42	0.44	0.29	0.08	0.36	0.24	0.31	0.33	0.29
Qwen3-Coder-30B-A3B-Instruct	0.15	0.31	0.04	0.28	0.20	0.06	0.12	0.18	0.13	0.14	0.15	0.10	0.13
Llama-3.3-70b-instruct	0.19	0.13	0.01	0.08	0.10	0.06	0.09	0.09	0.18	0.15	0.08	0.08	0.10
gemma3-12B	0.12	0.16	0.0	0.17	0.11	0.08	0.05	0.11	0.09	0.14	0.12	0.06	0.10
gemma3-27B	0.17	0.18	0.03	0.22	0.15	0.17	0.13	0.08	0.14	0.17	0.21	0.03	0.13
Llama3-8B	0.05	0.13	0.0	0.13	0.08	0.0	0.05	0.0	0.02	0.0	0.02	0.0	0.01
Llama3.1-8B	0.02	0.08	0.0	0.11	0.05	0.03	0.03	0.06	0.03	0.03	0.03	0.01	0.03
Mixtral-8x7B	0.01	0.08	0.0	0.12	0.05	0.01	0.01	0.01	0.0	0.01	0.01	0.01	0.01

Table 5: Complete Models' performance on algorithms of structure-oriented ideas.

LLMs	non-optimized algorithms				heuristic-optimized algorithms		
	BO	Recursion	Searching	Avg.	AS	IDAS	Avg.
DeepSeek-v3.2-Speciale	0.97	0.93	0.87	0.92	0.39	0.58	0.49
gemini-3-pro	0.97	0.91	0.78	0.89	0.34	0.25	0.30
grok4-vip	0.90	0.53	0.87	0.77	0.40	0.39	0.39
gpt-o3	1.0	0.89	0.80	0.90	0.32	0.47	0.39
Qwen3-235B	0.80	0.78	0.77	0.78	0.38	0.45	0.41
DeepSeek-R1	0.85	0.86	0.75	0.82	0.41	0.65	0.53
gemini-2.5-pro	0.84	0.87	0.69	0.80	0.24	0.26	0.25
Claude Opus	0.85	0.89	0.61	0.79	0.37	0.55	0.46
gpt-oss-120B	0.47	0.37	0.60	0.48	0.26	0.30	0.28
Qwen3-4B	0.82	0.86	0.41	0.70	0.22	0.40	0.31
Qwen3-8B	0.93	0.91	0.33	0.72	0.28	0.48	0.38
Qwen3-14B	0.87	0.94	0.53	0.78	0.28	0.54	0.41
Qwen3-32B	0.73	0.70	0.50	0.64	0.24	0.50	0.37
QwQ-32B	0.80	0.72	0.33	0.62	0.27	0.55	0.41
DeepSeek-R1-0528-Qwen3-8B	0.73	0.77	0.54	0.68	0.27	0.53	0.40
Distill-Qwen-32B	0.81	0.83	0.48	0.71	0.31	0.44	0.38
OpenThinker2-32B	0.82	0.81	0.51	0.71	0.28	0.52	0.40
Skyworker-OR1-32B	0.73	0.72	0.55	0.67	0.33	0.51	0.42
Qwen3-Coder-30B-A3B-Instruct	0.63	0.31	0.19	0.38	0.16	0.29	0.23
Llama-3.3-70b-instruct	0.47	0.69	0.20	0.45	0.10	0.32	0.21
gemma3-12B	0.37	0.63	0.16	0.38	0.08	0.31	0.19
gemma3-27B	0.50	0.77	0.23	0.50	0.16	0.29	0.23
Llama3-8B	0.17	0.06	0.10	0.11	0.02	0.06	0.04
Llama3.1-8B	0.13	0.33	0.08	0.18	0.06	0.13	0.10
Mixtral-8x7B	0.17	0.11	0.03	0.07	0.07	0.13	0.10




Figure 5: The results of models' parameters scaling.



Figure 6: Error type distributions across algorithmic reasoning taxonomies from Qwen3-4B to Qwen3-235B.




Figure 7: The co-occurrence heatmap of Qwen3-4B for low-entropy tokens proximal to indicator tokens (e.g., wait, but and so on) in global-optimized algorithmic reasoning. The blue segments along the x-axis denote necessary low-entropy tokens proximal to indicator tokens, whereas the orange segments correspond to regular high-entropy tokens appearing in the similar contextual positions.




Figure 8: The co-occurrence heatmap of Qwen3-8B for low-entropy tokens proximal to indicator tokens (e.g., wait, but and so on) in global-optimized algorithmic reasoning. The blue segments along the x-axis denote necessary low-entropy tokens proximal to indicator tokens, whereas the orange segments correspond to regular high-entropy tokens appearing in the similar contextual positions.




Figure 9: The co-occurrence heatmap of Qwen3-14B for low-entropy tokens proximal to indicator tokens (e.g., wait, but and so on) in global-optimized algorithmic reasoning. The blue segments along the x-axis denote necessary low-entropy tokens proximal to indicator tokens, whereas the orange segments correspond to regular high-entropy tokens appearing in the similar contextual positions.




Figure 10: The co-occurrence heatmap of Qwen3-32B for low-entropy tokens proximal to indicator tokens (e.g., wait, but and so on) in global-optimized algorithmic reasoning. The blue segments along the x-axis denote necessary low-entropy tokens proximal to indicator tokens, whereas the orange segments correspond to regular high-entropy tokens appearing in the similar contextual positions.




Figure 11: The co-occurrence heatmap of Qwen3-235B for low-entropy tokens proximal to indicator tokens (e.g., wait, but and so on) in global-optimized algorithmic reasoning. The blue segments along the x-axis denote necessary low-entropy tokens proximal to indicator tokens, whereas the orange segments correspond to regular high-entropy tokens appearing in the similar contextual positions.