

Why Do Open-Source LLMs Struggle with Data Analysis? A Systematic Empirical Study

Yuqi Zhu^{1,4}, Yi Zhong¹, Jintian Zhang^{1,4}, Ziheng Zhang³, Shuofei Qiao^{1,4},
Yujie Luo^{1,4}, Lun Du^{2,4}, Da Zheng^{2,4}, Ningyu Zhang^{1,4*}, Huajun Chen^{1,4*}

¹ Zhejiang University

² Ant Group

³ Independent Researcher

⁴ Zhejiang University - Ant Group Joint Laboratory of Knowledge Graph
{zhuyuqi,zhangningyu}@zju.edu.cn

Abstract

Large Language Models (LLMs) hold promise in automating data analysis tasks, yet open-source models face significant limitations in these kinds of reasoning-intensive scenarios. In this work, we investigate strategies to enhance the data analysis capabilities of open-source LLMs. By curating a seed dataset of diverse, realistic scenarios, we evaluate model behavior across three core dimensions: data understanding, code generation, and strategic planning. Our analysis reveals three key findings: (1) Strategic planning quality serves as the primary determinant of model performance; (2) Interaction design and task complexity significantly influence reasoning capabilities; (3) Data quality demonstrates a greater impact than diversity in achieving optimal performance. We leverage these insights to develop a data synthesis methodology, demonstrating significant improvements in open-source LLMs' analytical reasoning capabilities.

Code — <https://github.com/zjunlp/DataMind>

Extended version — <https://arxiv.org/abs/2506.19794>

1 Introduction

Data analysis is a complex, interactive process central to scientific discovery, business intelligence, and decision-making (Donoho 2017; Inala et al. 2024a). It requires models to understand natural language queries, interpret structured data, formulate hypotheses, generate executable code, and iteratively refine reasoning—often across multiple turns of interaction (Figure 1). As such, it demands tight integration of language understanding, logical reasoning, programming skills, and long-horizon planning, which together pose unique challenges beyond standard NLP tasks.

Large Language Models (LLMs) have shown promise in automating data analysis, with systems like DS-Agent (Guo et al. 2024), AutoML-Agent (Tirat, Jeong, and Hwang 2024), and Data Interpreter (Hong et al. 2024) demonstrating increasingly sophisticated behaviors. Domain-specific benchmarks such as DS-Bench (Jing et al. 2024), BLADE (Gu et al. 2024), QRData (Liu et al. 2024), and DiscoveryBench (Majumder et al. 2024) further enable targeted

* Corresponding Author.

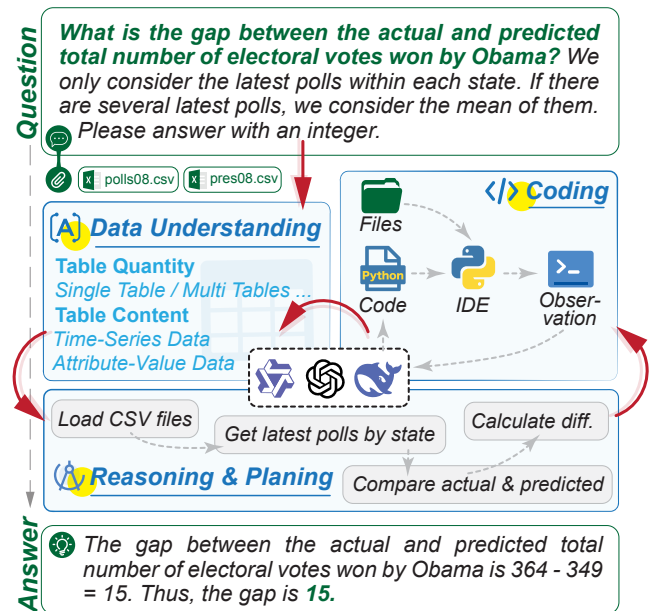


Figure 1: Core capabilities involved in data analysis tasks. We break down the process into three key components: data understanding, coding and planning.

evaluation. Yet, performance on these tasks remains dominated by large-scale, advanced models such as GPT-4 (Hurst et al. 2024) and DeepSeek-R1 (DeepSeek-AI et al. 2025), while open-source alternatives, especially smaller models, still struggle in real-world analytical settings.

This raises a key research question: *How can we effectively enhance open-source LLMs for complex, reasoning-intensive data analysis tasks?* Prior work has shown that fine-tuning on high-quality synthetic data can improve reasoning capabilities in domains like mathematics (Muenighoff et al. 2025; Ye et al. 2025) and code generation (Ahmad et al. 2025). However, for data analysis tasks that involve multi-step interactions, dynamic environments, and mixed goals, it remains unclear which properties of the training data, such as task difficulty, scenario diversity, or interaction structure, actually lead to better generalization.

In this paper, we take a capability-aware approach to understanding how to build more effective data analysis agents by analyzing the model capabilities involved in the analytical process. We decompose the task into three core dimensions: **Data Comprehension**, **Code Generation**, and **Strategic Planning**. Leveraging a curated seed dataset that encompasses diverse data analysis scenarios, we conduct targeted experiments and ablation studies to analyze the factors influencing model generalization and performance on complex analytical tasks. Here, our analysis reveals three key findings:

- The model’s planning ability emerges as a more critical determinant of success than its capabilities in data understanding and code generation. This underscores the importance of strategic foresight and structured reasoning in navigating complex analytical scenarios.
- Appropriate interaction turns, combined with the complexity of the data and suitable reasoning descriptions, can enhance the model’s reasoning capacity. However, performance gains vary across different data analysis tasks, suggesting that task-specific characteristics play a crucial role in shaping the model’s effectiveness.
- High-quality training data proves to be more critical than data diversity for achieving optimal performance in data analysis tasks. This emphasizes the necessity of curating datasets with precise and comprehensive annotations to ensure reliable and robust model outcomes.

To substantiate these insights, we propose a strategy-guided data synthesis framework that leverages empirical findings to inform key design choices, such as selecting informative interaction patterns and enriching data with concise reasoning traces, to enable more effective model learning. By fine-tuning on the resulting dataset, we demonstrate measurable improvements in performance, achieving results competitive with leading closed-source models.

2 Background of Data Analysis Agents

Data analysis tasks aim to derive actionable insights from data through systematic exploration and analysis (Inala et al. 2024b; Sun et al. 2024). In a typical workflow, analysts begin with specific questions about a dataset, then proceed through multiple analytical steps. These steps include data preprocessing and cleaning, hypothesis exploration, data transformations, and report generation. The process is inherently interactive - analysts work with structured tabular data, develop analysis code, interpret intermediate results, and iteratively refine their analysis before generating final reports with their findings.

To formally characterize the data analysis process for LLM agents, we define a parameterized analysis function f_θ that maps the input components to the analytical outputs:

$$f_\theta : (\mathcal{D}, \mathcal{Q}, \mathcal{T}) \rightarrow (\mathcal{S}, \mathcal{R}) \quad (1)$$

where \mathcal{D} represents the structured data, \mathcal{Q} specifies the analytical objective or query, and \mathcal{T} is the library of available analysis tools. f_θ models the behavior of an LLM-based agent that performs analysis by generating a sequence of intermediate analysis states $\mathcal{S} = \{s_t\}$, ultimately producing a final report \mathcal{R} to summarize the results.

3 Experimental Settings

Evaluation Protocol We adopt a capability-aware evaluation protocol aligned with our decomposition into Data Comprehension, Code Generation, and Strategic Planning. For the former two, we use *prompt-based evaluation*; for Strategic Planning, we employ *LoRA-based fine-tuning* (Hu et al. 2022) to target long-horizon reasoning.

Model Selection We evaluate a series of open-source model variants, including Qwen2.5-7B-Instruct, Qwen2.5-7B-Coder, Qwen2.5-7B-Instruct-1M, and R1-Distill-Qwen-7B, alongside strong API-based baselines like GPT-4o, DeepSeek-v3¹, and DeepSeek-R1. All models follow the ReAct framework (Yao et al. 2023) for multi-turn interaction, iteratively alternating between planning, code generation, and execution.

Data Collection and Curation To construct the training dataset, we collect samples from DAEval (Hu et al. 2024) (real-world CSVs from GitHub), DSbench (Jing et al. 2024) (ModelOff competition tasks), TableBench (Wu et al. 2025) (multi-domain tabular reasoning), WTQ (Pasupat and Liang 2015), and FetaQA (Nan et al. 2022). To enrich reasoning diversity, we generate additional synthetic samples all using DeepSeek-R1.

All training data are collected to ensure no overlap with our evaluation benchmarks, enabling a reliable out-of-distribution (OOD) assessment. Utilizing correctness-based filtering, we retain **6,443** distinct samples, encompassing a diverse array of analytical challenges. To ensure data quality, we apply a two-stage filtering process. First, we remove invalid samples based on several criteria: low-quality code implementations, such as those failing to utilize provided files or producing code with no meaningful return values; samples containing compilation errors; and entries that do not conform to format requirements. Following the automated filtering, we perform manual verification through sampling to further refine the dataset. The final corpus contains **5,613** high-quality instances used for fine-tuning.

Evaluation We evaluate on two comprehensive benchmarks: **DiscoveryBench** (Majumder et al. 2024), which includes 239 real-world tasks from its 264-task suite across six domains (e.g., sociology, engineering), and **QRData** (Liu et al. 2024), a benchmark designed for statistical and causal analysis comprising 411 questions paired with data sheets from textbooks, academic papers, and online resources.

Following Zeng et al. (2025), we use accuracy as the evaluation metric. Given that both predictions and references are in natural language, we employ GPT-4o-mini² for agreement scoring. Further details on dataset statistics and training configurations are provided in the extended version.

4 Core Capabilities for Data Analysis

Data analysis tasks present a unique challenge, requiring the integration of multiple capabilities. We identify three

¹gpt-4o-2024-08-06, deepseek-v3-0324

²gpt-4o-mini-2024-07-18

Model	QRData		DiscoveryBench	
	w/o Info	w Info	w/o Info	w Info
Qwen2.5-7B	6.57	7.54	0.42	1.26
Qwen2.5-14B	15.09	15.82	0.42	0.00

Table 1: Accuracy comparison of 7B and 14B models on QRData and DiscoveryBench (w/o and w/ table information).

core skills that determine model performance: **Data Comprehension**, the ability to understand and effectively utilize structured data; **Code Generation**, the skill of producing correct and efficient analytical code; and **Strategic Planning**, the capacity to decompose complex problems into manageable steps. All ablation experiments use comparable dataset sizes (via subsampling) to ensure fair within-condition comparisons. Slight variations across modules reflect differing experimental goals. This setup supports a systematic evaluation of the data characteristics that underlie effective analytical reasoning, thereby laying the foundation for our investigation.

4.1 Data Comprehension

To investigate whether data comprehension serves as a critical factor in enabling effective data analysis, we design a series of experiments to evaluate the model’s ability to reason over structured information. Specifically, the experiments focus on two key aspects: (1) whether explicitly providing structured context, such as tabular data, enhances the model’s reasoning accuracy, and (2) how the model performs when faced with increased complexity, including scenarios involving multiple sources of structured information, some of which may be irrelevant or distracting.

Tabular Information. To evaluate the impact of tabular data visibility, we compare two settings: without (**w/o Info**) and with (**w/ Info**) table information. In the **w/ Info** setting, we provide the necessary table information, such as column names, data types, and sample entries; in **w/o Info**, only the filename (e.g., `data.csv`) is provided. This ablation tests whether explicit access to table context improves reasoning. As shown in Table 1, adding table information helps the models in simpler tasks like QRData, showing slight improvements in performance. However, the gains are limited, indicating that the models already handle much of the reasoning without explicit table inputs. For the more complex DiscoveryBench, while the 7B model benefits from the inclusion of table information, the 14B model exhibits a slight drop in accuracy. This decline may be related to the increased input length, which could lead the 14B model to generate longer, less focused outputs. We hypothesize that this affects reasoning coherence, though further analysis is needed to confirm the effect.

Data Complexity. To evaluate the model’s ability to maintain focus amid distracting information, we introduce additional tables as irrelevant inputs, which act as semantic noise. This setup requires the model to reason over multi-

Model	QRData		DiscoveryBench	
	w/o Extra	w/ Extra	w/o Extra	w/ Extra
Qwen2.5-7B	37.96	34.55	5.44	4.18
Qwen2.5-14B	52.55	52.07	10.88	12.13

Table 2: Accuracy comparison of 7B and 14B models on QRData and DiscoveryBench (w/o and w/ extra data files).

Model	QRData	DiscoveryBench
Multi-Turn Setting		
Qwen2.5-7B-Instruct	39.71	14.64
Qwen2.5-14B-Instruct	53.53	24.27
Qwen2.5-32B-Instruct	57.18	27.62
Qwen2.5-7B-Coder	36.50	13.60
Qwen2.5-7B-Instruct-1M	39.17	15.48
R1-Distill-Qwen-7B	30.41	7.95
GPT-4o	59.85	28.03
DeepSeek-v3	65.21	36.82
Deepseek-R1	63.26	37.66

Table 3: Performance comparison across models in multi-turn settings (% accuracy).

ple tables, some of which are irrelevant, simulating scenarios with heightened complexity. Models that effectively filter out irrelevant data while focusing on task-relevant information are considered to exhibit stronger data understanding capabilities. For a fair evaluation, we remove task-specific background descriptions from the input, forcing the model to rely solely on the tabular inputs alone. As shown in Table 2, the inclusion of redundant data increases input complexity but does not lead to a significant decline in overall performance. The 7B model exhibits a modest decline, suggesting it is more sensitive to increased input noise. In contrast, the 14B model maintains stable performance, demonstrating stronger filtering capability and resilience in multi-source settings.

Discussion and Implications. The minimal performance gains from explicit table input suggest that data comprehension is not the primary bottleneck in data analysis. Even under input noise, models maintain stable performance and demonstrate robust data comprehension abilities—likely internalized during pretraining. This indicates that basic data understanding has already been internalized during pretraining, rendering additional context less beneficial.

4.2 Code Capability

To investigate the role of code generation in data analysis, we evaluate a diverse set of models with varying training objectives and architectures. Rather than treating code correctness as an end in itself, we examine how well models utilize code as part of a broader reasoning process to achieve task success.

Code Performance. We analyze how different models utilize code during problem solving. Table 3 provides the over-

Model	QRData	DiscoveryBench
Qwen2.5-7B-Instruct	34.64%	54.25%
Qwen2.5-14B-Instruct	29.94%	40.64%
Qwen2.5-32B-Instruct	20.73%	31.73%
Qwen2.5-7B-Coder	43.30%	50.98%
Qwen2.5-7B-Instruct-1M	31.44%	55.00%
R1-Distill-Qwen-7B	48.37%	60.00%

Table 4: Average code error rate of different models.




Figure 2: The distribution of error type.

all performance across models, while Table 4 summarizes the average code error rates observed across two datasets.

Our analysis reveals several key findings: (1) **Code specialization does not guarantee better performance:** Qwen2.5-7B-Coder does not demonstrate a clear advantage over general-purpose models. This suggests that code specialization alone may not translate to better performance in analytical tasks, due to limitations in instruction-following or reasoning generalization. (2) **Distillation can lead to functional hallucination:** R1-Distill-Qwen-7B, despite being distilled from a large reasoning model, performs poorly, often hallucinating file interpretations rather than generating executable code. (3) **Long-context capability does not imply efficient task execution:** When comparing Qwen2.5-7B-Instruct-1M and Qwen2.5-7B-Instruct under matched output length constraints, both models exhibit comparable coding capabilities; however, the latter demonstrated superior planning efficiency by completing tasks in fewer interaction rounds.

To further substantiate these findings, we manually sample 354 erroneous responses and categorized the errors using GPT-4o-mini. The categorization is based on the gap between the incorrect responses and the corresponding correct trajectories. As shown in Figure 2, only a small fraction of errors stem from syntactic or semantic code defects (e.g., invalid syntax). The majority of errors stem from higher-level reasoning failures, such as incorrect hypothesis formulation or premature termination. This indicates that *planning*, rather than *coding*, is more important.

Discussion and Implications. Our results indicate that while coding proficiency is necessary, it may not be the primary determinant of success in data analysis. Modern instruction-tuned models already possess sufficient code generation capabilities to handle typical analytical opera-

Turn Category	# Sample	QRData	DiscoveryBench
All	5613	48.66	15.00
Short	1034	47.68	23.85
Medium	3559	49.15	18.83
Long	1020	47.94	18.41
Medium + Short	4593	47.45	21.34
Medium + Long	4579	46.96	21.76

Table 5: Performance across different training data turn lengths (% accuracy).

tions. What truly distinguishes stronger agents may be their ability to *strategically deploy code*—to select appropriate operations, sequence them logically, and interpret outputs for iterative reasoning.

4.3 Strategic Planning

Data analysis is inherently a planning-intensive process that demands careful coordination of data access, transformation, and reasoning steps. Building on our earlier findings that strategic reasoning plays a decisive role in task success, we further investigate how the reasoning capabilities of LLMs shape their performance in complex analytical workflows. Given that instruction-tuned models (e.g., Qwen2.5-7B-Instruct) demonstrate more coherent planning behavior, we use them as the foundation for all subsequent experiments. Specifically, we systematically evaluate model performance along four key aspects: Interaction Turns, Reasoning Length, Task Complexity, and Problem Diversity.

Interaction Turns. To assess the impact of dialogue turn strategies on model performance, we categorized interactions into three primary turn lengths: **Short** (2-3 turns), **Medium** (4-5 turns), and **Long** (6+ turns). Additionally, we included a **Mixed** strategy that combines varying turn lengths to reflect more dynamic interaction scenarios.

Figure 3 reports the performance of Qwen2.5-7B and Qwen2.5-14B under different dialogue strategies. To ensure a fair comparison, both models were fine-tuned using the same dataset size across all strategies, which is 1020 here. The results reveal consistent trends across the two models: medium-length interactions generally achieve relatively better performance across both datasets, suggesting they provide an optimal balance between reasoning depth and focus. In contrast, short and long interaction strategies yield slightly lower yet relatively stable results. Notably, the mixed strategy consistently underperforms, likely because variable turn lengths disrupt the model’s ability to learn stable interaction patterns. Given the alignment in trends across model scales and the computational efficiency of the 7B variant, we select Qwen2.5-7B for all subsequent experiments. This enables deeper ablation studies while maintaining analytical rigor.

To better understand the impact of turn length, we examine the performance of various interaction strategies using the full collected dataset. The results are presented in Table 5, from which we derive the following observations:

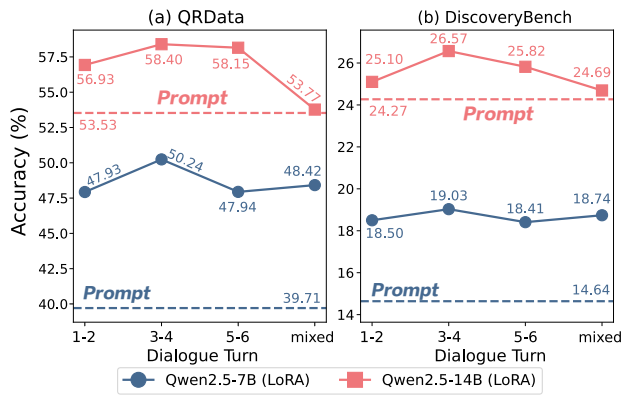


Figure 3: Impact of dialogue turn strategies across different Qwen model scales and training methods.

(i) **Turn length preferences are task-dependent.** The results reveal that Medium-length turns achieve a relatively higher performance on QRData, indicating their suitability for tasks requiring moderate reasoning depth. On DiscoveryBench, which features longer and more complex inputs, Short turns surprisingly outperform other strategies, potentially due to their ability to focus on concise and straightforward reasoning.

(ii) **Data quality outweighs quantity.** Our experiments reveal that increasing the amount of training data does not necessarily lead to better performance, even when using the same interaction turn strategies. In fact, medium-length turns trained on a smaller subset consistently outperform models trained on the full dataset. This finding highlights the importance of data quality and task relevance, suggesting that the effectiveness of fine-tuning is shaped by factors beyond data quantity alone.

(iii) **Mixed strategies require principled design.** While combining turn lengths may seem beneficial for diversity, the underperformance of mixed strategies suggests that unstructured variation introduces inconsistency, hindering the model’s ability to learn coherent reasoning patterns. Rather than improving flexibility, random mixing may confuse policy learning. Effective use of diverse interaction styles likely requires intentional design, such as curriculum scheduling or adaptive control, rather than uniform combination.

To ensure consistency and facilitate direct comparisons in subsequent experiments, we adopt the **Medium** turn strategy as the baseline for all follow-up evaluations.

Reasoning Length. To investigate whether longer reasoning chains from stronger models improve planning and task success, we augment training samples with intermediate `<think>` segments generated by DeepSeek-R1. These segments aim to capture intermediate reasoning steps that may scaffold better decision-making during multistep analysis.

We evaluate three settings: (1) **Original reasoning**, which uses the original training data without modification. These samples include reasoning content, but the reasoning is typically shorter and less explicit compared to the aug-

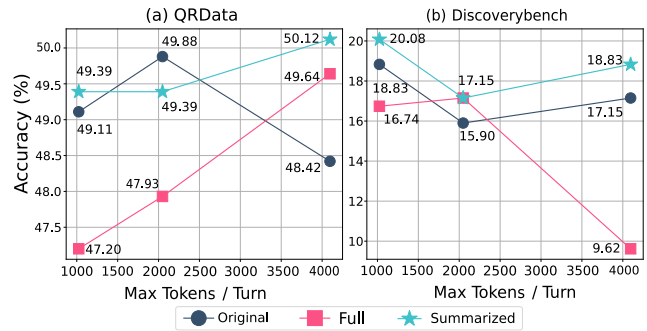


Figure 4: Impact of reasoning length on model performance across token budgets.

mented settings; (2) **Full reasoning**, which replaces the original reasoning with full `<think>` traces for intermediate turns (excluding the first and final turns); and (3) **Summarized reasoning**, which substitutes the original reasoning with concise summaries (generated by an LLM) of the full traces. The reasoning is inserted only in the middle turns to avoid hallucinations during initial grounding and to maintain brevity in the final answer generation. Experiments are conducted under varying per-turn token budgets (1024, 2048, and 4096) to simulate different interaction constraints. Figure 4 presents the results. We observe that:

(i) **Longer reasoning is not always better.** While longer reasoning chains might seem beneficial, the **Full** setting consistently underperforms the **Original** across most configurations, particularly in DiscoveryBench. The sharp decline in performance at the 4096-token level suggests that excessive, unfiltered reasoning can overwhelm the model, either by introducing redundancy, amplifying internal noise, or disrupting attention coherence. In contrast, the **Summarized** setting (though shorter) consistently matches or surpasses baseline accuracy. By distilling critical inferences and removing irrelevant reasoning steps, it reduces noise and enhances focus, leading to more reliable decision-making. These results indicate that reasoning effectiveness depends less on length and more on information relevance and logical coherence. Well-structured, goal-aligned reasoning often outperforms longer but unfocused alternatives.

(ii) **Token budgets exhibit diminishing returns.** Increasing the per-turn token budget can improve performance by enabling better integration of reasoning and task-relevant content, as seen in tasks like QRData. However, these gains are often limited, and in some cases, such as with DiscoveryBench, a larger token budget may even reduce performance by amplifying noise or irrelevant information. This suggests that effectiveness in reasoning-intensive tasks depends less on context size and more on information density and coherence. Simply allocating more tokens without improving content quality offers limited value.

Task Complexity. To assess how data difficulty affects model reasoning, we classify each example based on the performance of models with varying capacities. Specifically, a

Difficulty	QRData	DiscoveryBench
<i>Easy</i>	42.58	20.50
<i>Medium</i>	51.34	18.83
<i>Hard</i>	48.18	19.50
<i>Medium + Hard</i>	51.34	23.01

Table 6: Performance comparison across different task complexity (% accuracy).

task is labeled *easy* if it can be correctly solved by Qwen2.5-7B, *medium* if only Qwen2.5-14B can solve it, and *hard* if it requires DeepSeek-R1 to provide the correct answer. Here, each level contains 733 samples to ensure fair comparison.

We train models using data of varying difficulty levels, with the results summarized in Table 6. On QRData, performance generally improves with training data difficulty, with the *medium* setting yielding the best results. Combining *medium* and *hard* data achieves comparable or better performance across both datasets, indicating that exposure to more complex tasks enhances model generalization on structured analytical problems.

To understand this effect, we analyze the models’ dialogue turns and average response length, as illustrated in Figure 5. The results reveal that: **As training data difficulty increases, models shift from multi-turn, feedback-dependent interactions to generating comprehensive answers in fewer steps.** This suggests that exposure to harder tasks encourages models to internalize reasoning steps—reducing reliance on iterative refinement and increasing reasoning density per turn. The result is more efficient, self-contained decision-making.

Notably, training on a mix of *medium* and *hard* data yields the most compact behavior, with the lowest average interaction turns and shortest responses. However, on DiscoveryBench, models trained on *easy* data perform better. We attribute this to the need for fine-grained, multi-step computation, where a smaller per-step workload improves reliability. These findings highlight a key trade-off: while complex training fosters reasoning efficiency, simpler strategies may remain valuable for tasks requiring extended exploration.

Problem Diversity. We further investigate whether adjusting the diversity of question improves model performance. To this end, we annotate the dataset with semantic category labels using GPT-4o-mini and retain ten major domains after manual consolidation.

To quantify diversity, we employed a three-stage procedure to classify each question into distinct domains. The detailed methods and descriptions of the categories are detailed in the extended version of this paper. From the complete medium turn dataset, we select 2,220 examples under two settings: (1) a natural distribution that reflects the original domain frequencies, and (2) a balanced distribution that down-samples dominant domains while retaining all examples from underrepresented ones. Both settings preserve the overall difficulty distribution of the dataset. Results in Table 7 show minimal performance differences between the




Figure 5: Impact of training data difficulty on interaction patterns. (a) Average number of response rounds of the model. (b) Average output token length of the model.

Diversity	QRData	DiscoveryBench
Original Distribution	46.72	20.92
Balanced Sampling	45.00	21.76

Table 7: Performance under different sampling strategies (% accuracy).

two settings, indicating that domain diversity alone does not significantly influence model performance in this context. These findings suggest that the effectiveness of a model is not solely determined by the diversity of problems it encounters during training. Rather, the diversity and richness of reasoning strategies, such as the depth of reasoning processes and the complexity of logical steps, appear to play a more influential role in shaping a model’s performance.

What Makes Data Effective ? The efficacy of large language models in data analysis is fundamentally shaped by the quality of the data rather than its sheer diversity. High-quality, well-structured tasks that emphasize nuanced complexity and transparent reasoning pathways are instrumental in refining the models’ analytical capabilities.

5 Towards Effective Data Analysis Performance

While the previous section examined how different aspects of data analysis influence model performance, we now ask a practical question: how can these insights be leveraged to guide data collection and reuse in order to improve data analysis capabilities?

5.1 Strategy-Guided Data Synthesis

We design our data synthesis process in three systematically organized stages to construct a refined dataset that supports reasoning in data analysis. 1) **Prompt-Based Answer Generation.** We begin by leveraging prompt-based generation techniques to produce multiple candidate answers for each input query. This step ensures a diverse pool of responses, capturing a range of reasoning patterns and perspectives. 2)

Model	QRData	DiscoveryBench
API Models		
GPT-4o	59.85	28.03
DeepSeek-v3	65.21	36.82
Deepseek-R1	63.26	37.66
7B Models		
Qwen2.5-Instruct	39.71	14.64
Ours	53.77	22.59
14B Models		
Qwen2.5-Instruct	53.53	24.27
Ours	58.15	36.82

Table 8: Performance comparison of our method and baselines across model types and sizes (% accuracy).

Targeted Instance Selection. Next, we prioritize medium-length dialogues and examples of medium to high difficulty, as these have been shown to facilitate more stable and effective learning. Through this filtering process, we select instances that strike a balance between complexity and informativeness. 3) **Reasoning-Driven Data Enrichment.** Finally, we enrich each selected instance with a concise reasoning summary. This step enhances abstraction and generalization by explicitly capturing the underlying reasoning process, allowing models to better learn transferable insights. By following this three-stage synthesis process with format standardization, we construct a dataset comprising 2.8k instances. The resulting dataset serves as the foundation for Supervised Fine-Tuning (SFT) to optimize model performance. A detailed overview of the synthesis process and the training parameters can be found in the extended version of this paper.

5.2 Evaluation Results

Table 8 presents the evaluation results of our approach, where we fine-tune 7B and 14B models on the curated dataset. Notably, the fine-tuned 7B model demonstrates substantial performance improvements compared to its baseline, while the fine-tuned 14B model achieves results that are comparable to or surpass those of GPT-4o. These results suggest that even simple, insight-driven adjustments to training data can yield substantial improvements in models’ performance on complex analytical tasks. However, the performance gains appear to diminish as model scale increases, suggesting a potential saturation point. One possible explanation is that our strategy is constructed using Qwen2.5-7B, making the resulting training distribution better aligned with its inductive biases. While this alignment benefits smaller architectures, it may be less effective for larger models with more diverse representational capacities. And a key limitation of our approach lies in the dataset itself. While the curated dataset provides value, it falls short in addressing the diversity and complexity required for more challenging tasks, making high-quality data a persistent bottleneck. To overcome this, future efforts should focus on expanding the dataset to include richer, more representative samples from

real-world applications. Such improvements would not only better capture the variability and nuances of real-world scenarios but also enable further refinement of filtering strategies, enhancing scalability and generalizability across different model sizes and domains.

6 Related Work

LLM Agents. To adapt LLMs to complex reasoning tasks, recent work has explored three main approaches: prompt engineering, supervised fine-tuning (SFT), and reinforcement learning (RL) (Wang et al. 2024; Xi et al. 2025). Prompt-based methods (Yao et al. 2023; Wang et al. 2023) improve reasoning performance by reformulating inputs to better elicit the model’s latent capabilities. SFT-based methods (Yin et al. 2024; Zhu et al. 2025; Muennighoff et al. 2025) adapt LLMs to reasoning tasks by training on labeled data, aligning model behavior with desired task outputs. Notable examples include S1 (Muennighoff et al. 2025), which introduces "budget forcing" to control reasoning steps, and LIMO (Ye et al. 2025), which demonstrates that complex mathematical abilities can emerge from carefully curated examples. RL-based approaches (DeepSeek-AI et al. 2025; Wang et al. 2025; Chen et al. 2025) enhance reasoning by optimizing multi-step decision-making, further aligning the model with effective problem-solving strategies.

LLM agents for Data Analysis. Data science is an interdisciplinary field that focuses on extracting valuable insights from various data sources, with data analysis serving as a crucial component (Zhang et al. 2025a; Luo et al. 2025). Recently, researchers have proposed several specialized LLMs to enhance automated data analysis capabilities (Sun et al. 2024; Li et al. 2025). Data Copilot (Zhang et al. 2023) is a code-centric agent that uses pre-designed interfaces to automate massive data processing for financial analysis tasks. Data Interpreter (Hong et al. 2024) leverages hierarchical dependency graphs, enabling automatic task breakdown and code improvements. AutoKaggle (Li et al. 2024) employs a multi-agent system for end-to-end tasks with optional human input. To evaluate the ability of these agents, researchers have developed several comprehensive benchmarks, such as InfiAgent-DABench (Hu et al. 2024), DS-Bench (Jing et al. 2024), DA-code (Huang et al. 2024), and DataSciBench (Zhang et al. 2025b) for assessing the effectiveness of LLM-based data analysis solutions.

7 Conclusion and Future Work

In this work, we conduct a detailed investigation into the data efficiency challenges faced by open-source LLMs in data analysis tasks. By curating a dataset specifically for data analysis scenarios, we systematically assess how data structure, interaction patterns, and training strategies influence model performance. Our findings highlight that careful multi-turn data design and appropriately structured training data are critical for enhancing LLM reasoning in data analysis. In future work, we will focus on expanding our dataset and exploring broader model evaluations to further advance LLMs for data analysis.

Acknowledgements

We would like to express our sincere gratitude to the anonymous reviewers for their thoughtful and constructive feedback. This work was supported by the National Natural Science Foundation of China (No. 62576307, No. NSFCU23B2055, No. NSFCU19B2027), the Fundamental Research Funds for the Central Universities (226-2023-00138), Ningbo Natural Science Foundation (2024J020), Yongjiang Talent Introduction Programme (2021A-156-G), and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University. This work was supported by Ant Group and Zhejiang University - Ant Group Joint Laboratory of Knowledge Graph.

References

- Ahmad, W. U.; Narenthiran, S.; Majumdar, S.; Ficek, A.; Jain, S.; Huang, J.; Noroozi, V.; and Ginsburg, B. 2025. OpenCodeReasoning: Advancing Data Distillation for Competitive Coding. *CoRR*, abs/2504.01943.
- Chen, M.; Li, T.; Sun, H.; Zhou, Y.; Zhu, C.; Wang, H.; Pan, J. Z.; Zhang, W.; Chen, H.; Yang, F.; Zhou, Z.; and Chen, W. 2025. ReSearch: Learning to Reason with Search for LLMs via Reinforcement Learning. *CoRR*, abs/2503.19470.
- DeepSeek-AI; Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; Zhang, X.; Yu, X.; Wu, Y.; Wu, Z. F.; Gou, Z.; Shao, Z.; Li, Z.; Gao, Z.; Liu, A.; Xue, B.; Wang, B.; Wu, B.; Feng, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; Dai, D.; Chen, D.; Ji, D.; Li, E.; Lin, F.; Dai, F.; Luo, F.; Hao, G.; Chen, G.; Li, G.; Zhang, H.; Bao, H.; Xu, H.; Wang, H.; Ding, H.; Xin, H.; Gao, H.; Qu, H.; Li, H.; Guo, J.; Li, J.; Wang, J.; Chen, J.; Yuan, J.; Qiu, J.; Li, J.; Cai, J. L.; Ni, J.; Liang, J.; Chen, J.; Dong, K.; Hu, K.; Gao, K.; Guan, K.; Huang, K.; Yu, K.; Wang, L.; Zhang, L.; Zhao, L.; Wang, L.; Zhang, L.; Xu, L.; Xia, L.; Zhang, M.; Zhang, M.; Tang, M.; Li, M.; Wang, M.; Li, M.; Tian, N.; Huang, P.; Zhang, P.; Wang, Q.; Chen, Q.; Du, Q.; Ge, R.; Zhang, R.; Pan, R.; Wang, R.; Chen, R. J.; Jin, R. L.; Chen, R.; Lu, S.; Zhou, S.; Chen, S.; Ye, S.; Wang, S.; Yu, S.; Zhou, S.; Pan, S.; and Li, S. S. 2025. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *CoRR*, abs/2501.12948.
- Donoho, D. 2017. 50 years of data science. *Journal of Computational and Graphical Statistics*, 26(4): 745–766.
- Gu, K.; Shang, R.; Jiang, R.; Kuang, K.; Lin, R.; Lyu, D.; Mao, Y.; Pan, Y.; Wu, T.; Yu, J.; Zhang, Y.; Zhang, T. M.; Zhu, L.; Merrill, M. A.; Heer, J.; and Althoff, T. 2024. BLADE: Benchmarking Language Model Agents for Data-Driven Science. In *AI-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., Findings of the Association for Computational Linguistics: EMNLP 2024, Miami, Florida, USA, November 12-16, 2024*, 13936–13971. Association for Computational Linguistics.
- Guo, S.; Deng, C.; Wen, Y.; Chen, H.; Chang, Y.; and Wang, J. 2024. DS-Agent: Automated Data Science by Empowering Large Language Models with Case-Based Reasoning. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 16813–16848. PMLR.
- Hong, S.; Lin, Y.; Liu, B.; Liu, B.; Wu, B.; Li, D.; Chen, J.; Zhang, J.; Wang, J.; Zhang, L.; Zhang, L.; Yang, M.; Zhuge, M.; Guo, T.; Zhou, T.; Tao, W.; Wang, W.; Tang, X.; Lu, X.; Zheng, X.; Liang, X.; Fei, Y.; Cheng, Y.; Xu, Z.; and Wu, C. 2024. Data Interpreter: An LLM Agent For Data Science. arXiv:2402.18679.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Hu, X.; Zhao, Z.; Wei, S.; Chai, Z.; Ma, Q.; Wang, G.; Wang, X.; Su, J.; Xu, J.; Zhu, M.; Cheng, Y.; Yuan, J.; Li, J.; Kuang, K.; Yang, Y.; Yang, H.; and Wu, F. 2024. InfiAgent-DABench: Evaluating Agents on Data Analysis Tasks. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.
- Huang, Y.; Luo, J.; Yu, Y.; Zhang, Y.; Lei, F.; Wei, Y.; He, S.; Huang, L.; Liu, X.; Zhao, J.; and Liu, K. 2024. DA-Code: Agent Data Science Code Generation Benchmark for Large Language Models. In *AI-Onaizan, Y.; Bansal, M.; and Chen, Y., eds., Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024, Miami, FL, USA, November 12-16, 2024*, 13487–13521. Association for Computational Linguistics.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; Madry, A.; Baker-Whitcomb, A.; Beutel, A.; Borzunov, A.; Carney, A.; Chow, A.; Kirillov, A.; Nichol, A.; Paino, A.; Renzin, A.; Passos, A. T.; Kirillov, A.; and et al. 2024. GPT-4o System Card. arXiv:2410.21276.
- Inala, J. P.; Wang, C.; Drucker, S. M.; Ramos, G.; Dibia, V.; Riche, N.; Brown, D.; Marshall, D.; and Gao, J. 2024a. Data Analysis in the Era of Generative AI. *CoRR*, abs/2409.18475.
- Inala, J. P.; Wang, C.; Drucker, S. M.; Ramos, G.; Dibia, V.; Riche, N.; Brown, D.; Marshall, D.; and Gao, J. 2024b. Data Analysis in the Era of Generative AI. *CoRR*, abs/2409.18475.
- Jing, L.; Huang, Z.; Wang, X.; Yao, W.; Yu, W.; Ma, K.; Zhang, H.; Du, X.; and Yu, D. 2024. DSbench: How Far Are Data Science Agents to Becoming Data Science Experts? arXiv:2409.07703.
- Li, Y.; Moussa, H. N.; Chen, Z.; Chen, S.; Yu, B.; Xue, M.; Burns, B.; Chiu, T.-Y.; Dey, V.; Lu, Z.; et al. 2025. AutoSDT: Scaling Data-Driven Discovery Tasks Toward Open Co-Scientists. *arXiv preprint arXiv:2506.08140*.
- Li, Z.; Zang, Q.; Ma, D.; Guo, J.; Zheng, T.; Liu, M.; Niu, X.; Wang, Y.; Yang, J.; Liu, J.; Zhong, W.; Zhou, W.; Huang, W.; and Zhang, G. 2024. AutoKaggle: A Multi-Agent Framework for Autonomous Data Science Competitions. *CoRR*, abs/2410.20424.
- Liu, X.; Wu, Z.; Wu, X.; Lu, P.; Chang, K.; and Feng, Y. 2024. Are LLMs Capable of Data-based Statistical and Causal Reasoning? Benchmarking Advanced Quantitative Reasoning with Data. In *Ku, L.; Martins, A.; and Srikumar,*

- V., eds., *Findings of the Association for Computational Linguistics, ACL 2024, Bangkok, Thailand and virtual meeting, August 11-16, 2024*, 9215–9235. Association for Computational Linguistics.
- Luo, A.; Xian, X.; Du, J.; Tian, F.; Wang, G.; Zhong, M.; Zhao, S.; Bi, X.; Liu, Z.; Zhou, J.; et al. 2025. AssistedDS: Benchmarking How External Domain Knowledge Assists LLMs in Automated Data Science. *arXiv preprint arXiv:2506.13992*.
- Majumder, B. P.; Surana, H.; Agarwal, D.; Mishra, B. D.; Meena, A.; Prakhar, A.; Vora, T.; Khot, T.; Sabharwal, A.; and Clark, P. 2024. DiscoveryBench: Towards Data-Driven Discovery with Large Language Models. *CoRR*, abs/2407.01725.
- Muennighoff, N.; Yang, Z.; Shi, W.; Li, X. L.; Fei-Fei, L.; Hajishirzi, H.; Zettlemoyer, L.; Liang, P.; Candès, E. J.; and Hashimoto, T. 2025. sl: Simple test-time scaling. *CoRR*, abs/2501.19393.
- Nan, L.; Hsieh, C.; Mao, Z.; Lin, X. V.; Verma, N.; Zhang, R.; Kryscinski, W.; Schoelkopf, H.; Kong, R.; Tang, X.; Mutuma, M.; Rosand, B.; Trindade, I.; Bandaru, R.; Cunningham, J.; Xiong, C.; and Radev, D. R. 2022. FeTaQA: Free-form Table Question Answering. *Trans. Assoc. Comput. Linguistics*, 10: 35–49.
- Pasupat, P.; and Liang, P. 2015. Compositional Semantic Parsing on Semi-Structured Tables. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, 1470–1480. The Association for Computer Linguistics.
- Sun, M.; Han, R.; Jiang, B.; Qi, H.; Sun, D.; Yuan, Y.; and Huang, J. 2024. A Survey on Large Language Model-based Agents for Statistics and Data Science. *CoRR*, abs/2412.14222.
- Trirat, P.; Jeong, W.; and Hwang, S. J. 2024. AutoML-Agent: A Multi-Agent LLM Framework for Full-Pipeline AutoML. *CoRR*, abs/2410.02958.
- Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6): 186345.
- Wang, L.; Xu, W.; Lan, Y.; Hu, Z.; Lan, Y.; Lee, R. K.; and Lim, E. 2023. Plan-and-Solve Prompting: Improving Zero-Shot Chain-of-Thought Reasoning by Large Language Models. In Rogers, A.; Boyd-Graber, J. L.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, 2609–2634. Association for Computational Linguistics.
- Wang, Z.; Wang, K.; Wang, Q.; Zhang, P.; Li, L.; Yang, Z.; Yu, K.; Nguyen, M. N.; Liu, L.; Gottlieb, E.; Lam, M.; Lu, Y.; Cho, K.; Wu, J.; Fei-Fei, L.; Wang, L.; Choi, Y.; and Li, M. 2025. RAGEN: Understanding Self-Evolution in LLM Agents via Multi-Turn Reinforcement Learning. *arXiv:2504.20073*.
- Wu, X.; Yang, J.; Chai, L.; Zhang, G.; Liu, J.; Du, X.; Liang, D.; Shu, D.; Cheng, X.; Sun, T.; Li, T.; Li, Z.; and Niu, G. 2025. TableBench: A Comprehensive and Complex Benchmark for Table Question Answering. In Walsh, T.; Shah, J.; and Kolter, Z., eds., *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA*, 25497–25506. AAAI Press.
- Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; et al. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2): 121101.
- Yao, S.; Zhao, J.; Yu, D.; Du, N.; Shafran, I.; Narasimhan, K. R.; and Cao, Y. 2023. ReAct: Synergizing Reasoning and Acting in Language Models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Ye, Y.; Huang, Z.; Xiao, Y.; Chern, E.; Xia, S.; and Liu, P. 2025. LIMO: Less is More for Reasoning. *CoRR*, abs/2502.03387.
- Yin, D.; Brahman, F.; Ravichander, A.; Chandu, K. R.; Chang, K.; Choi, Y.; and Lin, B. Y. 2024. Agent Lumos: Unified and Modular Training for Open-Source Language Agents. In Ku, L.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, 12380–12403. Association for Computational Linguistics.
- Zeng, Q.; Jin, C.; Wang, X.; Zheng, Y.; and Li, Q. 2025. An Analyst-Inspector Framework for Evaluating Reproducibility of LLMs in Data Science. *CoRR*, abs/2502.16395.
- Zhang, D.; Zhoubian, S.; Cai, M.; Li, F.; Yang, L.; Wang, W.; Dong, T.; Hu, Z.; Tang, J.; and Yue, Y. 2025a. DatasciBench: An llm agent benchmark for data science. *arXiv preprint arXiv:2502.13897*.
- Zhang, D.; Zhoubian, S.; Cai, M.; Li, F.; Yang, L.; Wang, W.; Dong, T.; Hu, Z.; Tang, J.; and Yue, Y. 2025b. DataSciBench: An LLM Agent Benchmark for Data Science. *CoRR*, abs/2502.13897.
- Zhang, W.; Shen, Y.; Lu, W.; and Zhuang, Y. 2023. DataCopilot: Bridging Billions of Data and Humans with Autonomous Workflow. *arXiv preprint arXiv:2306.07209*.
- Zhu, Y.; Qiao, S.; Ou, Y.; Deng, S.; Lyu, S.; Shen, Y.; Liang, L.; Gu, J.; Chen, H.; and Zhang, N. 2025. KnowAgent: Knowledge-Augmented Planning for LLM-Based Agents. In Chiruzzo, L.; Ritter, A.; and Wang, L., eds., *Findings of the Association for Computational Linguistics: NAACL 2025, Albuquerque, New Mexico, USA, April 29 - May 4, 2025*, 3709–3732. Association for Computational Linguistics.