TARAS: Toward Enhancing LLM's Content-aware Capability for Table Reasoning

Anonymous ACL submission

Abstract

Current works for table reasoning with LLMs can be broadly categorized into two distinct approaches: based on text reasoning or based on code generation, which leverage natural language processing or programming paradigms respectively. However, these two methods both suffer from limited table content-aware capability, due to table's large scale reaching the context length limit of LLMs or structural bias in generated codes. This paper proposes TARAS, a table reasoning architecture that enhances LLM's table content-aware capability at any scale. Considering the various distributions of records related to various questions in the original table, a decompose-recombine algorithm is introduced to obtain a refined table by decomposing the original table into sub-tables and recombining the records related to question extracted from them. According to the characteristics of these tables, an adaptive strategy will be adopted to select different solvers to generate multiple candidate answers and assign priorities to them. Finally, a semantic-based voting mechanism is designed to fuse these answers to obtain the final response. The experiment shows that TARAS has effectively improved LLM's table content-aware capability, thus achieving state-of-the-art performance in various table reasoning tasks. Our code is available at https://anonymous.4open.science/r/TARAS-5CB9/.

1 Introduction

Table serves as a fundamental format for representing structured relational data. While current large language models (LLMs) excel at many text-based tasks without task-specific model structure or training data, relying only on designing input prompts (Brown et al., 2020; Wei et al., 2022; Wang et al., 2022; Zhou et al., 2022; Kojima et al., 2022; Li et al., 2022), they still face challenges in table reasoning because of the complex structured characteristics of tabular data. The existing LLM-based table reasoning works can be systematically classified into two principal paradigms: based on text reasoning (Chen, 2022; Ye et al., 2023) or based on code generation (Cheng et al., 2022; Zhang et al., 2023; Wang et al., 2024b; Nahid and Rafiei, 2024), employing linguistic processing or computational logic respectively.

In text reasoning, LLM is prompted to reason with question and table data using natural language, benefiting from LLM's semantic parsing capability constructed from training. (Ye et al., 2023) prompts LLM to return a subset of the given table containing only relevant records, and then answer the question based on this subset. Although this approach allows LLM to be aware of all table contents related to the question, thus improving the reliability of reasoning, the performance is limited by the context length of LLM, which restricts the application for larger tables. In code generation, LLM is prompted to generate code based on question and table schema to operate on the table, utilizing its inherent coding capability to adopt symbolic reasoning, such as SQL statements (Nahid and Rafiei, 2024) or python programs (Wang et al., 2024b). Although this approach achieves scale-agnostic adaptability for no need to access table contents, it relies heavily on LLM's inherent coding pattern to infer the relevant records from question and table schema, which may introduce structural bias into the generated codes (Jiang et al., 2024; Wang et al., 2024a). Consequently, the code may not accurately reflect the specific content of the table, leading to incorrect responses. Some works have explored the idea of integrating text reasoning and code generation. (Liu et al., 2023) has proposed mix self-consistency to determine the final response through voting; H-STAR (Abhyankar et al., 2024) prompts LLM to choose what strategy to use for analysis based on the task type. However, these works are just simple combinations, not substantially enhance LLM's table content-aware

capability.

In summary, previous works both suffer from **limited table content-aware capability**, which forces us to explore **how a model can be aware of table content at any scale**. We propose TARAS, the TAble Refinement with Adaptive Strategy for table reasoning. Considering the various distributions of records related to various questions in the original table, a decompose-recombine algorithm is designed to decompose the original table into several sub-tables, and recombine the relevant records extracted from them into a refined table, and an adaptive strategy is raised to generate candidate answers with LLM from sub-tables and refined table and assign different priorities to these answers. To reduce the inherent inconsistency of the response form generated by LLM, a semantic-based voting is conducted on all candidate answers. We evaluate TARAS using different LLMs as the backbone model on various table reasoning tasks. The experiment results indicate that TARAS outperforms all existing LLM-based baselines.

The major contributions of this paper are listed below:

- We propose an architecture for table reasoning, namely TARAS, that aims to enhance LLM's table content-aware capability at any scale.
- An algorithm based on decompose-recombine is introduced to obtain a refined table which LLM can handle effectively by decomposing the original table into sub-tables and recombining the relevant records extracted from them. Additionally, a process including adaptive solving and priority assigning is designed to get multiple candidate answers according to specific table types, they are then fused based on semantics to obtain the final response.
- Our method has achieved leading performance compared to other LLM-based methods in various table reasoning tasks. We further discussed the significant superiority of our method through more experiments.

2 Related Works

Table Reasoning with Pre-trained Model Researchers have proposed some pre-trained language models on table reasoning (Herzig et al., 2020; Liu et al., 2021; Zhao et al., 2022; Yin et al., 2020; Gu et al., 2022). They have rich knowledge learned from large-scale data by collecting or synthesizing them. However, these methods require a large amount of corpus related to the task, which results in a high expense. In addition, they may also lead to overfitting problems, reducing their generalization.

Table Reasoning with LLM Compared with pre-trained models, LLM performs well on various tasks with zero-shot prompts, demonstrating the strong generalization. Therefore, there have been many LLM-based table reasoning methods recently. The paradigm includes two types: code generation and text reasoning. For code generation, Binder (Cheng et al., 2022) prompts LLM to identify and solve the parts of the question which cannot be solved by the original program such as schema linking. TabSQLify (Nahid and Rafiei, 2024) prompts LLM to decompose table by SQL statements. ReAcTable (Zhang et al., 2023) introduces ReAct (Yao et al., 2022) into table reasoning. It prompts LLM using external tools to operate table, generating intermediate tables to gradually enhance data, thereby converting it into a more accessible format for an easier solution. Similarly, Chain-of-Table (Wang et al., 2024b) introduce CoT (Wei et al., 2022). The difference is that it defines a set of table operations, and then prompts LLM to generate a table reasoning chain that calls those operations. Compared with the above methods, Chain-of-Table partly reduces the bias caused by letting LLM generate codes directly. For text reasoning, Dater (Ye et al., 2023) prompts LLM to decompose table and question to solve complex tasks. Besides, H-STAR (Abhyankar et al., 2024) discussed the issue of integrating text reasoning and code generation methods. However, it just simply prompts LLM to choose what strategy to use for analysis, without delving into the strengths and the limitations of different strategies, thus lacking interpretability.

3 Methodology

3.1 Overview

We propose TARAS, an architecture for LLM-based table reasoning tasks. It is scale-free and can fully unleash the potential of LLM by enhancing its table content-aware capability. As illustrated in Figure 1, TARAS includes a table refiner that decomposes the original table into sub-tables and then recombines relevant records obtained by the extraction of sub-tables into a refined table; an adaptive solver that can adopt different strategies to obtain

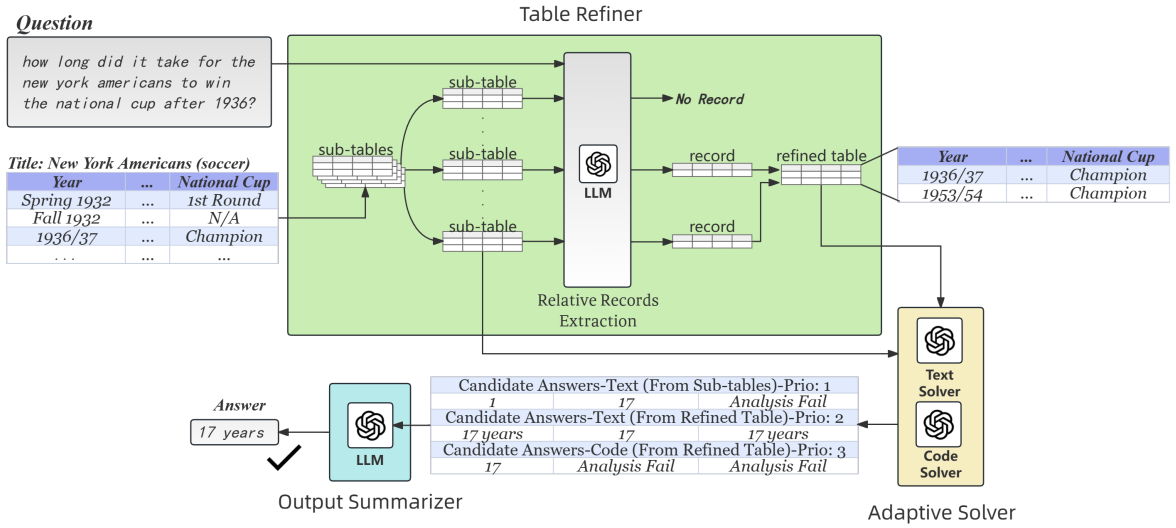

Figure 1: Overview of TARAS, which consists of three parts: (i) Table Refiner refines the original table into a refined table which has a higher density of relevant records; (ii) Adaptive Solver uses an adaptive strategy on different tables to generate multiple candidate answers and assign different priorities to them and (iii) Output Summarizer unifies semantically identical answers and vote to select the final answer.

multiple candidate answers; and an output summarizer for output processing, which can summarize all candidate answers and vote based on semantics to generate the final response.

3.2 Table Refiner

How to enhance the table content-aware capability of the model is the key for table reasoning task. Unlike other table-based tasks, table reasoning task requires the model to have a more specific understanding of the table down to the record level. Therefore, accurately locating relevant records is crucial. (Ye et al., 2023) used text reasoning to accomplish this task, prompting LLM to find all relevant records in the given table. However, it's difficult to search the entire table at once due to the context length limit of LLMs. Code generation method (Zhang et al., 2023; Wang et al., 2024b) prompts LLM to combine keywords from question and table schema to infer the relevant records, while the model does not actually understand the table data. Therefore, when the question is ambiguous, code generation raises errors.

We design an algorithm based on decompose-recombine to deal with it. The original table will be split into several sub-tables and a two-step sampling process for them is used to get the refined table. Algorithm 1 describes the overall process.

Table Split Line 1 to 2 of Algorithm 1 describes this step. The scale of sub-table should be considered to ensure that LLM can get enough information at once while avoiding long context lengths impairing the performance. We introduce a clipping mechanism to adjust the scale of each sub-table. Firstly, we divide the original table equally into three sub-tables. And then, if they are too small or too large, we will clip the scale to an appropriate range and re-divide.

Algorithm 1 Table Refinement Algorithm

Require: Table T, Question Q, Minimum Size of Sub-table l_{min}, Maximum Size of Sub-table l_{max}

Ensure: Table T'

1: $size$ = clip($\frac{T.length}{3}$, l_{min}, l_{max})
2: $subs$ = split_into_subtables(T, $size$)
3: Divide the $subs$ into 3 $clusters$.
4: **for** mid in each $clusters$ **do**
5: $prompt$ = gen_prompt(mid, Q)
6: $records$.append(LLM($prompt$))
7: **if** $records$ not null **then**
8: **for** each adj to the mid **do**
9: $prompt$ = gen_prompt(adj, Q)
10: $records$.append(LLM($prompt$))
11: **end for**
12: **end if**
13: **end for**
14: T' = $records$

Two-Step Sampling Process Lines 3 to 10 of Algorithm 1 describe this step. In many cases, the distribution of the relevant records in the table is relatively clustered, which means the high-cost action of searching all of these sub-tables can be avoided. Like "*What is the title of the next episode of a certain episode?*" or "*How long did it take for a certain team to win the championship again after a certain year?*", whose relevant records exist in adjacent sub-tables. To optimize the search process of sub-tables in above cases, we propose a two-step sampling process. **Step 1**: Divide all sub-tables into three clusters, and select the middle sub-table for each cluster. Lines 3 to 6 of Algorithm 1 describe this step. **Step 2**: Select two adjacent sub-tables from the same cluster if there are any relevant records in the selected sub-table. Lines 7 to 10 of Algorithm 1 describe this step.

Relevant Records Extraction LLM is prompted to judge whether the sub-table contains records related to the question, and if so, return their row numbers. After the extraction process is completed, all these rows are merged into a refined table for subsequent analysis.

3.3 Adaptive Solver

(Liu et al., 2023) compared the performance of using text reasoning or code generation to solve table reasoning tasks through many experiments. The results show that the accuracy of text reasoning performs better when the model can access the complete table. LLM is mainly trained on natural language data, which means its capability of code generation is inherently lower than text reasoning. However, large-scale tables in real-world scenarios limits the potential of LLM's text reasoning capability. Conversely, an obvious advantage of code generation is that it can answer a question under the circumstances of accessing only the schema of the table, allowing it to handle situations with a large table scale. However this advantage cannot fully cover the additional bias introduced by code generation methods due to the lack of awareness of table content. In summary, both these methods have their own strengths and limitations. We must carefully design the strategy that is flexible enough to face different situations.

An adaptive strategy is introduced to solve this problem. It can flexibly select different solvers to adapt on a sub-table or refined table to get multiple candidate answers and assign priorities to these answers. Algorithm 2 describes the overall process.

Algorithm 2 Adaptive Solving Algorithm

Require: Original Table T, New Table T', Set of Sub-tables T_s, Question Q, Minimum Size of Table l_{min}, Number of Calling Times n
Ensure: Set of Candidate Answers A

```
text_p = gen_text_prompt(T, Q)
text_res = LLM(text_p, n)
code_p = gen_code_prompt(T, Q)
code_res = LLM(code_p, n)
if T.length < l_min then
    assign_priority(code_res, 1)    ▷ value for reference only
    assign_priority(text_res, 2)
    A = text_res + code_res
else
    for each sub in T_s do
        sub_p = gen_text_prompt(sub, Q)
        sub_res = LLM(sub_p, n)
    end for
    assign_priority(text_res, 1)
    assign_priority(code_res, 2)
    assign_priority(sub_res, 3)
    A = sub_res + code_res + text_res
end if
```

Solvers Call LLM is prompted to act as different solvers. Text solver directly gives the answer or a "*can't be answered*" judgment, while code solver generates code to run on the original table to obtain the answer or an execution error. To reduce run-time errors in the generated code, the input and the output are defined in advance. To further enhance the diversity, self-consistency (Wang et al., 2022) that repeatedly calls the solvers multiple times is used to generate candidate answers.

Priority Assignment The adaptive strategy consists of three parts: **(1) Check the scale of original table**. The refinement is not necessary when the original table is small. In this case, the candidate answers from text reasoning will be assigned higher priorities than answer from code generation. **(2) The candidate answers from sub-tables should be given the highest priority**. The sub-tables sampled will not only be used to extract relevant records, but they will also be analyzed using text reasoning. That's because, first, if all relevant records exist in the same sub-table, this process will more likely obtain the correct answer. Second, since each sub-table contains different records, it's

unlikely that all candidate answers will be the same, so they will not cause significant interference to the subsequent analysis. In a word, a simple text reasoning will be conducted on all sub-tables, excluding failed processes and assigning the highest priority to the remaining candidate answers. **(3) Code is prior than text to analyze the refined table**. Regardless of the scale of the refined table, the candidate answers obtained by code generation will be assigned higher priority. Text reasoning is only suitable when all relevant records are accessible, while our table refinement algorithm faces the challenge of completeness when handling tables with long-tail structure or questions mentioning a large number of records. Therefore, the strategy tends to code generation. Moreover, previous LLM-based coding-dependent methods simply extracted the first few records of the table as the view provided to LLM. In contrast, our method provides LLM with a view that covers more relevant records, thereby improving the coding performance. In summary, all candidate answers are assigned different priorities through this step, which will affect the subsequent voting process.

3.4 Output Summarizer

Semantic Unification LLM prefers to add unnecessary context or explanation in response, leading to the inconsistency of the form. For example, for the question "*How old is Bob's father?*", from a semantic perspective, "*48*", "*48 years old*", "*Bob's father is 48 years old this year*" express the same meaning. But from a string comparison perspective, only one of them can be considered correct. Previous works mitigated this impact by designing complex string processing, while our method utilizes LLM to accomplish this task for its clear advantage in performing semantic-related tasks. Overall, in each process LLM is prompted to compare two candidate answers selected, and give a judgment on whether they are semantically identical. If yes, they will be unified into the same answer.

Prioritized Voting Voting will select the answer appears the most among all the candidate answers. If there are more than one answer appears the most, then the priority will be considered to make the final decision. It means the frequency is still preferred despite the possible low priority.

4 Experiments

4.1 Experimental Setup

Datasets We evaluate TARAS on four common table reasoning datasets: TabFact (Chen et al., 2019), WikiTQ (Pasupat and Liang, 2015), FeTaQA (Nan et al., 2022) and TableBench (Wu et al., 2024). TabFact is a table-based fact verification dataset, whose task is to judge whether a statement is true or not based on table data. We report results on its small test set. WikiTQ and FeTaQA are two table-based question answering datasets that focuses on answering questions through complex inference of the table data. FeTaQA is more difficult as it requires longer, freer responses. TableBench is a new comprehensive and complex benchmark that includes four major categories of table question answering tasks: numerical reasoning, fact checking, data analysis and visualization.

Baselines We report the LLM-based methods that have performed well on the above datasets, including Binder (Cheng et al., 2022), Dater (Ye et al., 2023), ReAcTable (Zhang et al., 2023), TabSQLify (Nahid and Rafiei, 2024), Chain-of-Table (Wang et al., 2024b) and H-STAR (Abhyankar et al., 2024). We also report some basic methods, including (a) Text End-to-End, which prompts LLM to directly give the answer, (b) Text CoT, which prompts LLM to think and give a reasoning chain before answering, (c) Code End-to-End, which prompts LLM to write code and (d) Code CoT, which prompts LLM to think and give a reasoning chain before writing code. Moreover, we also report some pre-trained methods including TaPas (Herzig et al., 2020), Tapex (Liu et al., 2021), ReasTAP (Zhao et al., 2022), OmniTab (Jiang et al., 2022), LEVER (Ni et al., 2023) and PASTA (Gu et al., 2022).

Metrics For TabFact, WikiTQ and TableBench, we directly compare the model response with the gold answer and judge whether they match. We don't use the Python-based WikiTQ evaluator (Pasupat and Liang, 2015), as our output summarizer takes this responsibility. We evaluate results on FeTaQA by calculating ROUGE (Lin, 2004) score of the model answer and the gold answer.

Experiment Settings The minimum size of the sub-table is set to 5, and the maximum size is set to 1.5 times the minimum size, rounded up. For the self-consistency setting, a solver will be called

Methods	Llama3.1-8b		Qwen2.5-7b		GPT-4o-mini	
	TabFact	WikiTQ	TabFact	WikiTQ	TabFact	WikiTQ
Text End-to-End	55.09	22.44	58.15	20.42	62.01	30.62
Text CoT	60.38	26.52	63.24	23.07	67.34	36.9
Code End-to-End	30.63	24.7	44.91	35.7	60.08	46.94
Code CoT	41.85	29.58	47.97	39.8	64.28	49.75
Binder	51.98	27.14	53.51	29.7	67.05	38.03
Dater	57.9	37.41	58.99	40.01	74.7	52
ReAcTable	-	37.64	-	40.42	-	57.09
TabSQLify	60.67	44.96	59.39	47.38	78.26	68.76
Chain-of-Table	63.83	-	62.35	-	85.13	-
H-STAR	70.1	51.06	64.38	51.24	89.38	74.95
TARAS	**75.69**	**58.49**	**73.22**	**59.46**	**89.82**	**78.5**

Table 1: Performance of various methods on TabFact and WikiTQ datasets. '-' indicates that we don't report the result because of the irreproducibility.

three times each process. In order to discuss the impact of different backbone models of TARAS, we evaluate it based on Llama3.1-8b (Dubey et al., 2024), Qwen2.5-7b (Yang et al., 2024) and GPT-4o-mini (Achiam et al., 2023). We don't report some results for the authors have not open-sourced the prompt and the code they used on the dataset.

Methods	TabFact	WikiTQ
TaPas	83.9	48.8
Tapex	85.9	57.5
ReasTAP	86.2	58.6
OmniTab	-	62.8
LEVER	-	65.8
PASTA	**90.6**	-
TARAS with Llama3.1-8b	75.69	58.49
TARAS with Qwen2.5-7b	73.22	59.46
TARAS with GPT-4o-mini	89.82	**78.5**

Table 2: Performance of pre-trained methods and TARAS with different backbone models on TabFact and WikiTQ datasets.

4.2 Experiment Results

Table 1 shows some of the results on comparison with previous LLM-based methods. Due to space limitations, we have the additional results in Appendix A.5. In summary, TARAS outperforms all existing LLM-based methods with different datasets and backbone models. It is worth mentioning that previous LLM-based methods sampled from the train set to construct a few-shot prompt, while we do not use any examples, which reduces the cost of generalizing our method to other table reasoning tasks and leaves room for further improvement.

Table 2 shows the comparison with previous pre-trained methods.

Table 3 shows results on comparison with previous LLM-based methods on TableBench dataset. TableBench includes four types of table reasoning tasks: numerical reasoning, fact checking, data analysis and visualization. Considering the absolute advantage of coding-dependent methods in visualization task, we exclude this task. Task in TableBench raises a greater challenge to LLM's table content-aware capability, as it contains more complex data and questions. Nevertheless, TARAS still achieved a comprehensive lead. The results indicate that TARAS effectively improves LLM's table content-aware capability, making it capable of handling more challenging table reasoning tasks.

4.3 Ablation Study

To evaluate the effectiveness of each module in TARAS, we conduct an ablation study on the WikiTQ dataset. We remove the table refiner, instead we always select the first sub-table. We remove the adaptive solver and only use the text solver or the code solver to generate candidate answers. Moreover, we shuffle the candidate answers to eliminate the influence of the priority. We remove the output summarizer, instead choosing the first candidate answer as the result.

Table 4 shows the specific contribution of each

Methods	Llama3.1-8b			Qwen2.5-7b			GPT-4o-mini		
	NR	FC	DA	NR	FC	DA	NR	FC	DA
Text End-to-End	32.49	52.08	10.79	37.78	50	11.08	48.36	64.58	22.45
Text CoT	40.81	60.42	15.74	42.32	58.3	18.66	60.71	66.67	26.53
Code End-to-End	36.02	56.25	17.49	46.1	53.13	20.99	52.9	65.62	35.28
Code CoT	45.09	64.58	19.24	52.64	65.63	24.2	61.46	67.71	36.15
TabSQLify	53.15	61.46	18.95	54.16	62.5	25.07	66.25	70.83	35.57
H-STAR	56.17	69.79	21.87	57.43	68.75	27.11	69.52	79.17	39.07
TARAS	**58.19**	**71.88**	**25.36**	**59.95**	**70.83**	**30.03**	**72.54**	**83.33**	**40.23**

Table 3: Performance of various methods on TableBench dataset. 'NR', 'FC' and 'DA' in table respectively means 'NumericalReasoning', 'FactChecking' and 'DataAnalysis' task in TableBench.

	Llama3.1-8b	GPT-4o-mini
TARAS	**58.49**	**78.5**
w/o Table Refiner	49.06	71.59
w/o Adaptive Solver(ALL Text Solver)	48.73	68.76
w/o Adaptive Solver(ALL Code Solver)	44.66	67.4
w/o Output Summarizer	44.34	56.4

Table 4: Performance of TARAS removing (i) Table Refiner, (ii) Adaptive Solver and (iii) Output Summarizer on WikiTQ dataset.

module to the TARAS's performance. Overall, removing any module will result in a performance loss, emphasizing the importance of all modules in TARAS.

Methods	Table Scale(Num. of Tokens)		
	<2k	2k-4k	>4k
Binder	33.59	26.65	5.13
Dater	48.65	28.53	13.68
ReAcTable	46.33	30.09	20.51
TabSQLify	52.7	37.93	30.77
H-STAR	56.95	45.45	40.17
TARAS	**65.25**	**52.04**	**46.15**

Table 5: Performance of various methods on different scales of tables from WikiTQ dataset, using Llama3.1-8b as backbone model.

4.4 Discussion

(Chen, 2022) points out that when the number of table rows exceeds 30, the performance bottleneck that a LLM can handle has been reached. To evaluate the impact of the table scale, we not only consider the number of rows but also the lengths. Specifically, we classify tables in the WikiTQ dataset into three parts by the number of tokens: tables with less than 2000 tokens, tables with more than 4000 tokens and tables with tokens from 2000 to 4000.

The results are in Table 5. All methods show a decrease in performance when the table scale increases. Relatively, TARAS controls the magnitude, thus ensuring a consistent performance as much as possible. This is attributed to our extraction algorithm based on decompose-recombine and the adaptive priority assignment strategy successfully reducing the impact of table scale on the performance.

Failed Step	Proportion
Two-Step Sampling	11%
Records Extracting	**35%**
Adaptive Solving	32%
Output Summarizing	22%

Table 6: Analysis of error types in 100 failure cases on the WikiTQ dataset, using Llama3.1-8b as backbone model.

We have collected 100 failure cases and recorded for each case which step that the error first appeared in. The results shown in Table 6 indicate that errors occur most frequently in table refiner, especially in

the step of relevant records extraction. It means that sometimes although LLM can access to the relevant records, they still cannot find out the relevance. Besides, even if we use adaptive strategy to answer the question, there are still many errors that occur in this step. That is to say, LLM can discover the relevance between question and records, but it does not know how to utilize it. In summary, enhancing the content-aware capability of LLM for table reasoning still has a long way to go. The meaning of "content-aware" lies not only in enabling LLM to access the data, but also in understanding and using them to solve the question.

Figure 2: Two cases on the WikiTQ dataset showing how modules in TARAS work together to finish the task.

Further, we select two cases for analysis that can fully demonstrate the effectiveness of TARAS. As illustrated in Figure 2. In case 1, the question mentions "*consecutive*", which is difficult to express in code. Although code solvers all fail, with table refiner's help, the text solver has obtained the correct result. While in case 2, conversely, table refiner fails to find all the relevant records for some reason, resulting in the text solver generating wrong answers. However, the performance of code solver has not been affected, thus maintaining the final response. These two cases have demonstrated that TARAS is not just a simple integration of previous methods, but an organic whole formed by the close connections between these modules. We have discussed more details about how text solvers and code solvers work together in A.2.

In the last, we discuss the efficiency of TARAS. We have calculated the running costs, including average input tokens(AIT), average output tokens(AOT), average execute time(AET) and average LLM call times(ACT), of various LLM-based methods. The results are shown in Table 7.

Methods	AIT	AOT	AET	ACT
Binder	91.34k	2.03k	125.07	44.76
Dater	7.5k	1.99k	43.91	4.62
ReAcTable	20.11k	1.77k	61.86	10.87
TabSQLify	**5.44k**	**1.04k**	**28.45**	**3.8**
H-STAR	12.29k	1.45k	50.07	6.89
TARAS	<u>6.44k</u>	3.11k	57.15	14.34

Table 7: Running costs statistics of various methods on the WikiTQ dataset, using Llama3.1-8b as backbone model. <u>underline</u> means the second best performance.

We construct zero-shot prompts, which reduces the count of input tokens. In addition, we implement TARAS using concurrent programming to accelerate the process. However, TARAS involves frequent LLM calls, resulting in significantly higher output tokens count and LLM calls. We emphasize that it is a sacrifice made to ensure the quality and stability of the response. But how to find a trade-off between performance and efficiency is the topic we will study next.

5 Conclusion

In this paper we propose TARAS, an architecture for table reasoning that aims to enhance LLM's content-aware capability at any scale. It includes a decompose-recombine algorithm to refine the original table; an adaptive strategy to obtain various answers and assign priorities to them; and a semantic-based voting process to generate the final response. Experiment results have shown that TARAS outperforms all existing LLM-based methods in various table reasoning tasks. Further analysis highlights the significance of our method.

Limitations

Our method has not been experimented on more table reasoning tasks such as Text-to-SQL, table-based prediction and tabular data analysis, etc., which imposes a limit on the generalization. In addition, the datasets used are constructed based on Wikipedia, the tabular data is simply organized, clear, easy to understand and not too large. However, tables in real-world scenarios raise more challenges including extremely large scale, noisy records and complex organization forms. How to address these challenges remains an unresolved

issue.

Moreover, our method further increases the complexity compared to previous LLM-based methods, involving repeated calls and interactions between multiple LLMs, which may lead to a possible decrease in efficiency. How to find a trade-off between performance and efficiency is the topic we will study next.

References

Nikhil Abhyankar, Vivek Gupta, Dan Roth, and Chandan K. Reddy. 2024. H-star: Llm-driven hybrid sql-text adaptive reasoning on tables. *ArXiv*, abs/2407.05952.

OpenAI Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haim ing Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, and 260 others. 2023. Gpt-4 technical report.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Wenhu Chen. 2022. Large language models are few(1)-shot table reasoners. *ArXiv*, abs/2210.06710.

Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. 2019. Tabfact: A large-scale dataset for table-based fact verification. *arXiv preprint arXiv:1909.02164*.

Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir R. Radev, Marilyn Ostendorf, Luke S. Zettlemoyer, Noah A. Smith, and Tao Yu. 2022. Binding language models in symbolic languages. *ArXiv*, abs/2210.02875.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony S. Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 510 others. 2024. The llama 3 herd of models. *ArXiv*, abs/2407.21783.

Zihui Gu, Ju Fan, Nan Tang, Preslav Nakov, Xiaoman Zhao, and Xiaoyong Du. 2022. Pasta: table-operations aware fact verification via sentence-table cloze pre-training. *arXiv preprint arXiv:2211.02816*.

Jonathan Herzig, Paweł Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Martin Eisenschlos. 2020. Tapas: Weakly supervised table parsing via pre-training. *arXiv preprint arXiv:2004.02349*.

Juyong Jiang, Fan Wang, Jiasi Shen, Sungju Kim, and Sunghun Kim. 2024. A survey on large language models for code generation. *ArXiv*, abs/2406.00515.

Zhengbao Jiang, Yi Mao, Pengcheng He, Graham Neubig, and Weizhu Chen. 2022. Omnitab: Pre-training with natural and synthetic data for few-shot table-based question answering. *arXiv preprint arXiv:2207.03637*.

Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *ArXiv*, abs/2205.11916.

Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, B. Chen, Jian-Guang Lou, and Weizhu Chen. 2022. Making language models better reasoners with step-aware verifier. In *Annual Meeting of the Association for Computational Linguistics*.

Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Annual Meeting of the Association for Computational Linguistics*.

Qian Liu, Bei Chen, Jiaqi Guo, Morteza Ziyadi, Zeqi Lin, Weizhu Chen, and Jian-Guang Lou. 2021. Tapex: Table pre-training via learning a neural sql executor. *arXiv preprint arXiv:2107.07653*.

Tianyang Liu, Fei Wang, and Muhao Chen. 2023. Rethinking tabular data understanding with large language models. In *North American Chapter of the Association for Computational Linguistics*.

Md Mahadi Hasan Nahid and Davood Rafiei. 2024. Tabsqlify: Enhancing reasoning capabilities of llms through table decomposition. *ArXiv*, abs/2404.10150.

Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryściński, Hailey Schoelkopf, Riley Kong, Xiangru Tang, and 1 others. 2022. Fetaqa: Free-form table question answering. *Transactions of the Association for Computational Linguistics*, 10:35–49.

Ansong Ni, Srini Iyer, Dragomir R. Radev, Ves Stoyanov, Wen tau Yih, Sida I. Wang, and Xi Victoria Lin. 2023. Lever: Learning to verify language-to-code generation with execution. *ArXiv*, abs/2302.08468.

Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. *arXiv preprint arXiv:1508.00305*.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.

Zhijie Wang, Zijie Zhou, Da Song, Yuheng Huang, Shengmai Chen, Lei Ma, and Tianyi Zhang. 2024a. Where do large language models fail when generating code? *arXiv preprint arXiv:2406.08731*.

Zilong Wang, Hao Zhang, Chun-Liang Li, Julian Martin Eisenschlos, Vincent Perot, Zifeng Wang, Lesly Miculicich, Yasuhisa Fujii, Jingbo Shang, Chen-Yu Lee, and Tomas Pfister. 2024b. Chain-of-table: Evolving tables in the reasoning chain for table understanding. *ArXiv*, abs/2401.04398.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Xianjie Wu, Jian Yang, Linzheng Chai, Ge Zhang, Jiaheng Liu, Xinrun Du, Di Liang, Daixin Shu, Xianfu Cheng, Tianzhen Sun, Guanglin Niu, Tongliang Li, and Zhoujun Li. 2024. Tablebench: A comprehensive and complex benchmark for table question answering. *ArXiv*, abs/2408.09174.

Qwen An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxin Yang, Jingren Zhou, Junyang Lin, and 25 others. 2024. Qwen2.5 technical report. *ArXiv*, abs/2412.15115.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. React: Synergizing reasoning and acting in language models. *ArXiv*, abs/2210.03629.

Yunhu Ye, Binyuan Hui, Min Yang, Binhua Li, Fei Huang, and Yongbin Li. 2023. Large language models are versatile decomposers: Decomposing evidence and questions for table-based reasoning. *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*.

Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. Tabert: Pretraining for joint understanding of textual and tabular data. *arXiv preprint arXiv:2005.08314*.

Yunjia Zhang, Jordan Henkel, Avrilia Floratou, Joyce Cahoon, Shaleen Deep, and Jignesh M. Patel. 2023. Reactable: Enhancing react for table question answering. *Proc. VLDB Endow.*, 17:1981–1994.

Yilun Zhao, Linyong Nan, Zhenting Qi, Rui Zhang, and Dragomir Radev. 2022. Reastap: Injecting table reasoning skills during pre-training via synthetic reasoning examples. *arXiv preprint arXiv:2210.12374*.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and 1 others. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

A Supplementary Experiments

A.1 Effectiveness of Table Refiner

Table refiner refines the original table into a table containing more relevant records and less redundant records through the algorithm based on decompose-recombine. Figure 3 compares the average scale of the original table and the refined table. We exclude all small tables that do not require refinement. It can be seen that table refiner significantly reduces the number of records in the refined table - from an average of 27.69 records to 6.16 records for the WikiTQ dataset, and from an average of 14.27 records to 4.25 records for the TabFact dataset. It demonstrates how our table refiner effectively reduces the redundant information unrelated to the question and allows the LLM to focus more on valuable information for inference.

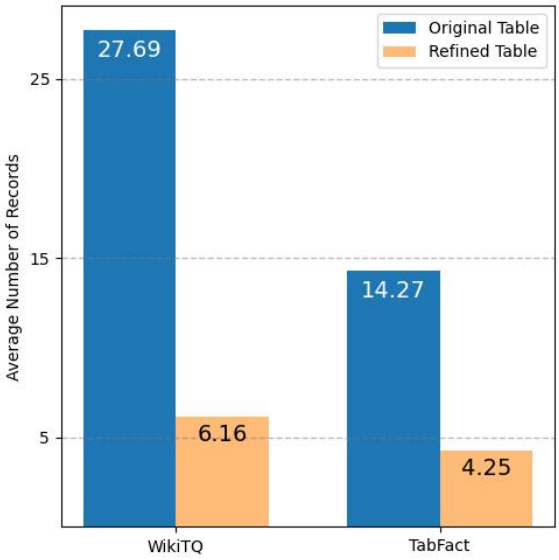

Figure 3: Comparison of the scale of original table/refined table. We exclude all the small tables as they do not require refinement.

A.2 Effectiveness of Adaptive Solver

We calculate the proportion of correct answers appearing only in the text solver/only in the code solver/in both/in neither. Table 8 shows the results. It can be seen that (a) the most common situation is that both the text solver and the code solver can obtain the correct answer, accounting for 31.66% and 61.32%. It indicates that there is actually no essential difference between these two methods, but how

	Llama3.1-8b	GPT-4o-mini
in Text Solver	23.27%	13.42%
in Code Solver	18.34%	11.95%
in Both	**31.66%**	**61.32%**
in Neither	26.73%	13.31%

Table 8: The proportion of correct answer in text solver/code solver/both of them/neither of them on the WikiTQ dataset.

we make choices based on the actual situations is important. In addition, after using a stronger LLM, the proportion of this situation has reached an absolute advantage, which once again proves the viewpoint that as the capability of the LLM improves, it can obtain the correct answer using any analysis method. (b) The correct answer appears more frequently in the text solvers than in the code solvers. There are two reasons for this. First, text solvers can keep a stable performance on a small table, while code solvers are likely to generate code with errors on any table. Second, when analyzing sub-tables, we only use the text solver. In some cases, all relevant records may exist in the same sub-table, so the text solver can directly obtain the correct answer, which cannot be achieved by code solvers.

A.3 Benefit of Using Self-Consistency

	Llama3.1-8b	GPT-4o-mini
TARAS	**58.49**	**78.5**
w/o SC	45.49	67.29

Table 9: Performance of TARAS with/without Self-Consistency on WikiTQ dataset.

Table 9 shows the impact of with/without self-consistency on the performance of TARAS. We use self-consistency only when analyzing the refined table but not the sub-tables, so the weight of candidate answers obtained from sub-tables has been weakened. In other words, our strategy ensures that this part of candidate answers have the highest priority while reducing their negative impact from the randomness of these answers to the result. It's worth noticing that after using a stronger LLM as the backbone model, the performance loss without self-consistency decreased from 22.23% to 14.28%. It can be explained by the fact that as the capability of the LLM improves, it can obtain the correct answer using any analysis method.

A.4 Benefit of Using Semantic Unification

	Llama3.1-8b	GPT-4o-mini
TARAS	**58.49**	**78.5**
w/o Unification	38.26	67.52

Table 10: Performance of TARAS with/without Semantic Unification on WikiTQ dataset.

Table 10 shows the impact of with/without semantic unification on the performance of TARAS. After removing this step, there was an astonishing and significant decline in the result. We emphasize that it DOES NOT demonstrate the importance of unification itself. If there were no such process, the candidate answers will only be voted on with simple string matching, thus leading to our adaptive solving method losing its expected consistency in the result. In other words, the performance degradation without unification is actually a continuation of the impact of our adaptive solving method on TARAS's performance, which proves the importance of this module from the side.

A.5 Evaluation on FeTaQA dataset

Table 11 shows the result of some LLM-based methods on the FeTaQA dataset with different backbone models. Unlike other table reasoning tasks, FeTaQA requires the model to provide a descriptive statement as the answer rather than directly giving the result. And for evaluation, FeTaQA calculates the matching degree between the model statement and the gold statement to assess the quality of the answer. To address this requirement, we have slightly modified the prompts to make the LLM tend to generate longer and more detailed answers.

It's worth mentioning that using only simple text reasoning + CoT can achieve a performance that is close to or even better than LLM-based methods. We attribute it to the evaluation rule of the FeTaQA. ROUGE score measures the completeness of the information related to the gold answer in the model answer. In fact, we observed that when the model generates a longer answer, it's easier to obtain a higher ROUGE score, regardless of whether the answer is correct. In addition, LLM excels in using descriptive statements to answer a question, which allows it to generate a response that has a similar style to the gold answer through a simple prompt.

Methods	Llama3.1-8b			Qwen2.5-7b			GPT-4o-mini		
	1	2	L	1	2	L	1	2	L
Text End-to-End	0.42	0.21	0.37	0.44	0.22	0.38	0.56	0.34	0.48
Text CoT	**0.48**	0.24	0.42	0.46	0.24	0.41	0.58	0.35	**0.52**
Code End-to-End	0.27	0.13	0.25	0.29	0.13	0.25	0.36	0.18	0.33
Code CoT	0.28	0.13	0.25	0.31	0.15	0.28	0.38	0.19	0.34
TabSQLify	0.41	0.2	0.37	0.37	0.16	0.33	0.49	0.35	0.44
H-STAR	0.45	0.23	0.4	0.46	0.25	0.43	0.54	0.44	0.5
TARAS	**0.48**	**0.25**	**0.45**	**0.49**	**0.27**	**0.44**	**0.63**	**0.46**	**0.52**

Table 11: Performance of various methods on FeTaQA dataset. '1' and '2' means calculating the score using ROUGE-n where n = 1 or 2. 'L' means calculating the score using ROUGE-L.