

AILS-NTUA at SemEval-2025 Task 8: Language-to-Code prompting and Error Fixing for Tabular Question Answering

Andreas Evangelatos, Giorgos Filandrianos, Maria Lymperaïou,
Athanasios Voulodimos, Giorgos Stamou

National Technical University of Athens

andrasevangelatos@mail.ntua.gr, {geofila, marialymp}@ails.ece.ntua.gr,
thanosv@mail.ntua.gr, gstam@cs.ntua.gr

Abstract

In this paper, we present our submission to SemEval-2025 Task 8: Question Answering over Tabular Data. This task, evaluated on the DataBench dataset, assesses Large Language Models' (LLMs) ability to answer natural language questions over structured data while addressing topic diversity and table size limitations in previous benchmarks. We propose a system that employs effective LLM prompting to translate natural language queries into executable code, enabling accurate responses, error correction, and interpretability. Our approach ranks first in both subtasks of the competition in the proprietary model category, significantly outperforming the organizer's baseline.

1 Introduction

The integration of Large Language Models (LLMs) into question-answering (QA) systems over tabular data has garnered significant attention in recent research, due to their capability to translate natural language into structured queries (Fang et al., 2024). First attempts in language-to-query conversion rely on fine-tuning pre-trained language models, so that they acquire tabular understanding (Liu et al., 2022). The vast contextual knowledge of LLMs, as well as their exposure on multiple modalities, including tabular comprehension, enables them to implicitly translate natural language in structured queries. As tabular QA arises as an emergent LLM ability (Chen, 2023), more and more prompting-based methods showcase improvements (Hegselmann et al., 2022; Nam et al., 2024).

Evaluating LLMs as tabular reasoners was explored with the DataBench dataset (Osés Grijalba et al., 2024), comprising 65 datasets from various domains, accompanied by manually crafted questions in natural language. Answers fall into different categories, such as boolean, categorical, numerical or list. Experiments utilizing in-context learning and code-based prompting reveal that while

LLMs show promise, there is significant room for improvement, particularly in smaller scales.

In this paper, we explore a wide range of prompting strategies for Large LLMs to identify optimal methods for converting natural language queries into code and effectively interacting with tabular data. Specifically, we propose a system for generating executable Python functions from natural language queries through LLM prompting. Our approach achieves accuracy scores of 85.63% and 87.93% in the DataBench and DataBench Lite tasks respectively, ranking first in both tasks in the proprietary models category (Osés-Grijalba et al., 2025). After human evaluation, the accuracy increases to 89.85% and 88.89% in DataBench and DataBench Lite respectively. In our system, LLMs are explicitly instructed to articulate their reasoning process, improving the reliability of their outputs and facilitating interpretability.

Our code is available on GitHub¹.

2 Background

2.1 Task description

The data used for this task are sourced from DataBench (Osés Grijalba et al., 2024). The task consists of queries in the form of (T, Q) where T is a structured table containing a set of c columns $\{col_1, \dots, col_c\}$ and a set of r rows $\{row_1, \dots, row_r\}$. Each row consists of c values $\{row_{i,1}, \dots, row_{i,c}\}$, which correspond to the entries in the table for the respective columns. These values can be of various data types commonly found in real-world datasets in English, such as numbers, categorical values, dates, and text. Additionally, some values may be missing. The expected answer to the query is a value a of type boolean, category, number, list[category] or list[number]. Answers are computed based on a small subset of the columns

¹<https://github.com/andrewevag/tabularqa>

and may either be values directly present in T or include statistics computed on these values.

Two subtasks are proposed with regards to the size of the table T of the input query:

Subtask I: DataBench QA – A table of any size is provided along with a question in natural language.

Subtask II: DataBench Lite QA – The same task is performed using a sampled version of the tables, where a maximum of 20 rows is retained.

2.2 Related work

LLMs excel in various tasks but struggle with tabular data, especially when key information is dispersed across large tables and complex queries (Ye et al., 2023). This necessitates shifting from direct LLM reasoning to generating intermediate representations in SQL or Python (Zhang et al., 2023b). For instance, Zhang et al. (2023a) serialize table schemas, enabling GPT-3 to iteratively refine SQL queries by correcting syntax and compilation errors. Similarly, Plan-of-SQLs (POS) (Giang et al., 2024) decomposes complex queries, reducing reliance on advanced Text-to-SQL models. PAL (Gao et al., 2023) generates intermediate programmatic steps, outsourcing execution to a Python interpreter. Lei et al. (2023) integrate multiple LLMs into a toolkit, leveraging Chain of Thought (CoT) (Wei et al., 2023) and Program of Thoughts (PoT) (Chen et al., 2022) prompting to enhance SQL-based reasoning. TabLLM (Zha et al., 2023) further extends comprehension by jointly training LLMs on tables and text. Building on this, Ye et al. (2023) improve table-based reasoning by segmenting large tables into relevant sub-tables and decomposing complex queries. Unlike intermediate representation-based methods, Wu and Hou (2025) use in-context learning with retrieval-augmented generation (RAG) for direct tabular query answering. Agarwal et al. (2025) enhance in-context learning by integrating tabular data and queries into a unified hybrid graph. Zhu et al. (2024) employ a step-wise pipeline to distill knowledge, enabling smaller models for discrete reasoning over tabular data.

3 System Overview

Our system mainly performs text-to-Python code conversions to queries via prompting. This allows for a lightweight and task-agnostic implementation that employs LLMs without further training. Specifically, the system consists of two modules:

the Main Module and the Error-Fixing Module. An overview of the architecture provided in Figure 1.

3.1 Main Module

This module is responsible for translating the query (T, Q) into executable Python code. First, a prompt is generated by applying a predefined template to the input (T, Q) . This prompt is then provided to an LLM, which completes a Python function P that takes the table T as a DataFrame input and, when executed, answers the question Q . The generated code is executed by a Python interpreter, which calls the function with the table T and outputs the predicted answer \hat{a} . If execution completes without errors, \hat{a} is considered the system’s final prediction. Otherwise, the Error-Fixing Module is invoked.

Prompt The prompt template f is used to structure the input of the LLM based on a (T, Q) pair, incorporating the following fields:

1. Task Description (TD): A comment describing the objective of completing a Python function to answer Q , followed by the function header.

2. Column Description (CD): A multiline comment listing the column indices, names, data types, and example values in CSV format.

3. Sample Rows (SR): A comment containing the first n rows of the table. Columns with values exceeding n_c characters are truncated, and values are right-aligned with spaces (Rajkumar et al., 2022).

4. Columns Used and Answer Type (CUAT): Comments specifying the columns required to answer Q , their data types, and the expected answer type.

Since CUAT is unknown for a new (T, Q) pair, generating P involves two steps. First, the LLM is prompted with (T, Q) using an incomplete version of f , denoted as f_{inc} , which contains TD, CD, and SR, to infer CUAT. Then, the inferred CUAT is incorporated into f , forming a complete prompt with all four fields.

In both steps, we leverage in-context learning by incorporating k manually crafted examples $\{(T_i, Q_i, P_i)\}_{i=1}^k$, where each (T_i, Q_i) pair represents a table and its corresponding query, and P_i is a program that processes T_i to answer Q_i . Each example is structured using the same template f , (CUAT is extracted by the P_i), and is appended with the corresponding P_i . The complete prompt for the first step is:

$$p_1 = (f(T_1, Q_1) || P_1) || \dots || (f(T_k, Q_k) || P_k)$$

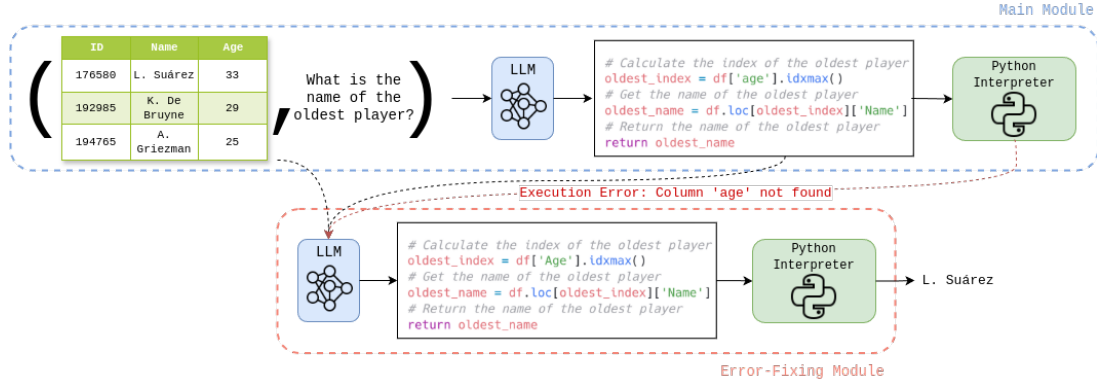


Figure 1: Overview of the architecture of our system.

$||f_{\text{inc}}(T, Q)$ while for the second: $p_2 = (f(T_1, Q_1)||P_1)|| \dots ||(f(T_k, Q_k)||P_k)||f(T, Q)$, where p_1, p_2 are the prompts used for Steps 1 and 2, respectively, and the CUAT field in $f(T, Q)$ of p_2 is computed by $g(p_1)$, where g is the output of LLM given a prompt. Further details on prompt construction are provided in the Appendix A.

Example Programs Furthermore, to enhance program generation, we incorporate a CoT strategy. This involves inserting comments describing intermediate reasoning steps, interleaved with the Python statements implementing these steps. Our approach aligns with methods used in PAL (Gao et al., 2023) and PoT (Chen et al., 2022), aiming to leverage LLMs’ reasoning capabilities while mitigating computational inaccuracies by offloading calculations to the Python interpreter. Meaningful variable names are also included in example programs, as they can further facilitate reasoning by clarifying intent.

Upon observing the example programs, the LLM generates a completion of the Python function P to answer Q , incorporating its reasoning in comments. The generated program is then executed by the Python interpreter, which outputs the predicted answer \hat{a} . If execution timeouts or fails, both the generated code and a description of the error are passed to the Error-Fixing Module.

3.2 Error-Fixing Module

This module utilizes an LLM prompted to correct execution errors in the function P . The prompt includes the following components:

1. **Task Description:** A description of the objective, instructing the model to fix the error.
2. **Faulty Program:** The erroneous code presented in the format $f(T, Q)||P$.

3. **Code Error:** A description message of the execution error.

4. **Function Header:** The corrected function signature that the model will complete.

Once the LLM generates the corrected function, it is executed in the Python interpreter to produce the answer \hat{a} . If execution completes successfully without timing out or failing, the system’s final prediction is \hat{a} . Otherwise, it retries until the maximum attempts and outputs Error if all fail.

4 Experimental Setup

Dataset The dataset used for model development is publicly available on Hugging Face². It contains 1,308 questions across 65 tables. We derive our exemplars from the predefined train split, comprising 988 questions from 49 tables. The columns required to answer each question, their data types, and the answer’s data type are precomputed by the task organizers and included in the dataset. Our annotations focus solely on the expected code and reasoning steps generated by the LLM. Model evaluation is conducted on the dev split, which includes 320 questions from 16 tables, as well as on the test set, which consists of 522 questions from 15 tables and serves as the basis for the final ranking.

Evaluation Metric Model performance is assessed using a relaxed accuracy measure, which accounts for predictions with equivalent semantic meaning, e.g. the prediction "Yes" is considered correct if the gold answer is True. Numerical answers are truncated to two decimal places before being compared against the gold answer. The evaluation metric is provided by the task organizers through the databench_eval³ Python package.

²<https://huggingface.co/datasets/cardiffnlp/databench>

³https://github.com/jorses/databench_eval

Baseline The baseline system employs a model to convert the input query into a Python function that takes the table as a Pandas DataFrame and returns the answer to the question. The generated function is then executed by the Python interpreter. The prompt includes detailed task instructions along with a simple example of a completed function. Additionally, it provides the names of the columns in the input DataFrame as a list. The model utilized is stable-code-3b-GGUF⁴.

Models and Hyperparameters We evaluate the performance of multiple LLMs as components of our system. In all experiments, the same LLM is used for both the Main and Error-Fixing Modules. The models include open-source variants of Llama (8B, 70B, and 405B), proprietary models like Claude 3.5 Sonnet, and code-generation models such as Qwen 2.5-Coder 7B. Details on model versions are in Appendix D.

For code generation, we select a temperature of $\text{temp}=0$ and $\text{top_p}=0.9$ in all experiments, except for the Error-Fixing Module for smaller models ($\leq 70\text{B}$ parameters), where we set $\text{temp}=1$ and allow up to three invocation attempts. The Main Module processes five rows from the table, while the Error-Fixing Module processes ten rows for both DataBench and DataBench Lite. The prompt includes nine exemplar cases, ensuring coverage of all possible answer data types in DataBench.

Resources Experiments involving Claude 3.5 Sonnet and Llama 3.1 Instruct 405B are conducted using Amazon Bedrock, while experiments with the remaining models are performed using Ollama on 4 NVIDIA A10G GPUs.

5 Results

Rankings Table 1 presents accuracy scores on dev and test set for our system, considering LLMs of different sizes. We rank 1st in both subtasks in the proprietary models track when employing Claude 3.5 Sonnet, which demonstrates substantial improvements over the baseline, surpassing it by $\sim 60\%$ in both subtasks. Moreover, our system ranks 3rd in DataBench in the open-source category using Llama 3.1 Instruct 405B. Our system significantly outperforms the baseline even when employing smaller models ($\leq 8\text{B}$ parameters). The accuracy scores of both smaller models rank just

below the 2nd and 3rd positions in DataBench and DataBench Lite respectively, in the smaller model ranking. However, due to submission limitations in the official competition, smaller models’ predictions were not submitted and are therefore not reflected in the final competition rankings.

Large vs. Small Models Our findings indicate that our system’s performance is heavily influenced by the overall capabilities of the LLM used. Larger models achieve up to 20% higher accuracy in both subtasks, which is expected given their superior code generation capabilities. This is evident even when compared with a smaller model that is specifically trained for code generation tasks.

Information in the Prompt We conduct an *ablation study* on the different fields included in the prompt used in the Main Module to assess their impact on performance. Specifically, we evaluate the accuracy of the Main Module using Llama 3.1 Instruct 405B as the underlying LLM, incrementally increasing the amount of information provided in the prompt until it matches the final system configuration. The results are reported in Table 2, including the following configurations:

1. **Simple Example + TD + CD + SR:** The prompt includes the fields TD, CD, and SR, along with a simple example of a completed function that calculates the number of rows in a DataFrame, as used in the baseline.

2. **Few-Shot (FS):** The prompt includes examples of completed functions using single-line Pandas statements. The LLM is instructed to generate function completions in a similar format.

3. **FS CoT:** The prompt includes examples of completed functions with the reasoning steps in comments. The LLM is instructed to generate responses following the same approach.

4. **FS TD + CD + SR:** The prompt includes TD, CD, and SR, along with examples of completed functions using single-line Pandas statements. The LLM is instructed to generate single-line function completions.

5. **FS CoT + TD + CD + SR:** The prompt includes all fields of the Main Module, excluding CUAT. The exemplars contain reasoning steps as comments, and the LLM is prompted to generate responses following the same approach.

We observe that the sole inclusion of column and row descriptions provides minimal improvement in model accuracy on the test set, while performance decreases on the dev set. However,

⁴<https://huggingface.co/TheBloke/stable-code-3b-GGUF>

Models	DataBench		DataBench Lite	
	Dev Set Accuracy	Test Set Accuracy	Dev Set Accuracy	Test Set Accuracy
Claude 3.5 Sonnet	91.87	85.63 (89.85)	87.81	87.93 (88.89)
Llama 3.1 Instruct 405B	91.56	83.33 (87.16)	90.00	77.78 (78.54)
Llama 3.3 70B	89.06	79.50	86.88	83.33
Llama 3.1 8B	79.06	65.13	74.38	67.82
Qwen2.5-Coder 7B	77.5	65.33	79.69	68.20
Baseline	-	26.00	-	27.00

Table 1: Accuracy scores (%) of our system across different LLMs on both subtasks. **Bold** denotes the submitted predictions and in the parenthesis are the official accuracy scores of those on DataBench after human reviewing.

Prompt	DataBench		DataBench Lite	
	Dev Set Accuracy	Test Set Accuracy	Dev Set Accuracy	Test Set Accuracy
TD	48.44	47.70	48.44	48.08
TD + CD	46.25	47.89	49.06	44.69
TD + CD + SR	40.00	48.08	41.25	48.28
Simple Example + TD + CD + SR	53.45	58.43	57.5	60.92
Few Shot	71.88	64.18	73.12	65.33
Few Shot CoT	75.31	65.71	79.69	67.43
Few Shot TD + CD + SR	85.94	78.93	76.05	74.52
Few Shot CoT + TD + CD + SR	91.25	79.50	80.31	76.05
Main Module	89.38	82.57	79.06	71.65

Table 2: Accuracy scores (%) for Llama 3.1 Instruct 405B with ablated prompts.

Model	DataBench - Test Set Accuracy			DataBench Lite - Test Set Accuracy		
	Main-Module	Full System	#Errors Fixed	Main-Module	Full System	#Errors Fixed
Claude 3.5 Sonnet	83.52	85.63	11/16	85.82	87.93	11/16
Llama 3.1 Instruct 405B	82.57	83.33	4/10	71.65	77.78	32/55
Llama 3.3 70B	77.97	79.5	8/22	81.61	83.33	9/20
Llama 3.1 8B	62.84	65.13	12/42	64.94	67.82	15/50
Qwen2.5-Coder 7B	63.22	65.33	11/49	65.90	68.20	12/44

Table 3: Accuracy scores (%) for the test set of Main Module and Full System, and number of errors fixed by the Error-Fixing Module for all models. **Bold** highlights best accuracy results across models.

their significance becomes apparent when multiple exemplars illustrating the task are included in the prompt. In the FS setting, model performance significantly improves across all datasets and subtasks; Notably, the mere incorporation of a single task demonstration increases accuracy by more than 10% across all datasets and subtasks. Finally, allowing the LLM to explicitly express its reasoning—detailing the information extracted from the table, the expected outcome, and the intermediate steps to derive the answer—leads to a consistent improvement in performance across all tasks.

Significance of the Error-Fixing Module We present the test set accuracy scores with and without the involvement of the Error-Fixing Module, along with the number of correctly fixed errors—resulting in the gold label—in Table 3. The Error-Fixing Module provides an average accuracy increase of 2.4% across all models and subtasks.

Larger models, such as Claude 3.5 Sonnet and Llama 3.1 Instruct 405B, produce fewer execution errors and successfully correct more than 40% of

them across all tasks. In contrast, smaller models generate a higher number of errors and demonstrate limited error-correction capabilities, successfully resolving only 30% of their execution errors at most. While further improvements to the Error-Fixing Module could substantially enhance the accuracy of smaller models, their impact on larger models would likely be minimal, as these models are inherently less prone to generating code that results in execution errors.

6 Conclusion

In this work, we present our system for SemEval-2025 Task 8: Question-Answering Over Tabular Data, which translates natural language questions into executable Python functions using only LLM prompting. Our approach decomposes the task into smaller steps, explicitly demonstrating its reasoning process, and incorporates a self-correction mechanism to handle execution errors. As a result, our system ranks first in both subtasks within the proprietary model category.

References

- Ankush Agarwal, Ganesh S, and Chaitanya Devaguptapu. 2025. [Hybrid graphs for table-and-text based question answering using llms](#). *Preprint*, arXiv:2501.17767.
- Wenhu Chen. 2023. [Large language models are few\(1\)-shot table reasoners](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1120–1130, Dubrovnik, Croatia. Association for Computational Linguistics.
- 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*.
- Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. 2023. [Binding language models in symbolic languages](#). *Preprint*, arXiv:2210.02875.
- Xi Fang, Weijie Xu, Fiona Anting Tan, Jiani Zhang, Ziqing Hu, Yanjun Qi, Scott Nickleach, Diego Socolinsky, Srinivasan Sengamedu, and Christos Faloutsos. 2024. [Large language models \(llms\) on tabular data: Prediction, generation, and understanding – a survey](#). *Preprint*, arXiv:2402.17944.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Pal: Program-aided language models](#). *Preprint*, arXiv:2211.10435.
- Carlos Gemmell and Jeffrey Dalton. 2023. [Generate, transform, answer: Question specific tool synthesis for tabular data](#). *Preprint*, arXiv:2303.10138.
- Giang, Nguyen, Ivan Brugere, Shubham Sharma, Sanjay Kariyappa, Anh Totti Nguyen, and Freddy Lecue. 2024. [Interpretable llm-based table question answering](#). *Preprint*, arXiv:2412.12386.
- Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David A. Sontag. 2022. [Tabllm: Few-shot classification of tabular data with large language models](#). *ArXiv*, abs/2210.10723.
- Fangyu Lei, Tongxu Luo, Pengqi Yang, Weihao Liu, Hanwen Liu, Jiahe Lei, Yiming Huang, Yifan Wei, Shizhu He, Jun Zhao, and Kang Liu. 2023. [Tableqakit: A comprehensive and practical toolkit for table-based question answering](#). *Preprint*, arXiv:2310.15075.
- Qian Liu, Bei Chen, Jiaqi Guo, Morteza Ziyadi, Zeqi Lin, Weizhu Chen, and Jian-Guang Lou. 2022. [Tapex: Table pre-training via learning a neural sql executor](#). *Preprint*, arXiv:2107.07653.
- Jaehyun Nam, Woomin Song, Seong Hyeon Park, Jihoon Tack, Sukmin Yun, Jaehyung Kim, Kyu Hwan Oh, and Jinwoo Shin. 2024. [Tabular transfer learning via prompting llms](#). *ArXiv*, abs/2408.11063.
- Jorge Osés Grijalba, L. Alfonso Ureña-López, Eugenio Martínez Cámara, and Jose Camacho-Collados. 2024. [Question answering over tabular data with DataBench: A large-scale empirical evaluation of LLMs](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13471–13488, Torino, Italia. ELRA and ICCL.
- Jorge Osés-Grijalba, Luis Alfonso Ureña-López, Eugenio Martínez Cámara, and Jose Camacho-Collados. 2025. SemEval-2025 task 8: Question answering over tabular data. In *Proceedings of the 19th International Workshop on Semantic Evaluation (SemEval-2025)*, Vienna, Austria. Association for Computational Linguistics.
- Nitarshan Rajkumar, Raymond Li, and Dzmitry Bahdanau. 2022. [Evaluating the text-to-sql capabilities of large language models](#). *Preprint*, arXiv:2204.00498.
- Zhijie Wang, Zijie Zhou, Da Song, Yuheng Huang, Shengmai Chen, Lei Ma, and Tianyi Zhang. 2025. [Towards understanding the characteristics of code generation errors made by large language models](#). *Preprint*, arXiv:2406.08731.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Jie Wu and Mengshu Hou. 2025. [An efficient retrieval-based method for tabular prediction with LLM](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 9917–9925, Abu Dhabi, UAE. Association for Computational Linguistics.
- 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](#). In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23*, page 174–184, New York, NY, USA. Association for Computing Machinery.
- Liangyu Zha, Junlin Zhou, Liyao Li, Rui Wang, Qingyi Huang, Saisai Yang, Jing Yuan, Changbao Su, Xi-ang Li, Aofeng Su, Tao Zhang, Chen Zhou, Kaizhe Shou, Miao Wang, Wufang Zhu, Guoshan Lu, Chao Ye, Yali Ye, Wentao Ye, Yiming Zhang, Xinglong Deng, Jie Xu, Haobo Wang, Gang Chen, and Junbo Zhao. 2023. [Tablegpt: Towards unifying tables, nature language and commands into one gpt](#). *Preprint*, arXiv:2307.08674.
- Hengyuan Zhang, Peng Chang, and Zongcheng Ji. 2023a. [Bridging the gap: Deciphering tabular data using large language model](#). *Preprint*, arXiv:2308.11891.

Yunjia Zhang, Jordan Henkel, Avrilia Floratou, Joyce Cahoon, Shaleen Deep, and Jignesh M. Patel. 2023b. [Reactable: Enhancing react for table question answering](#). *Preprint*, arXiv:2310.00815.

Fengbin Zhu, Ziyang Liu, Fuli Feng, Chao Wang, Moxin Li, and Tat Seng Chua. 2024. [Tat-llm: A specialized language model for discrete reasoning over financial tabular and textual data](#). In *Proceedings of the 5th ACM International Conference on AI in Finance, ICAIF '24*, page 310–318, New York, NY, USA. Association for Computing Machinery.

A Prompts

A.1 Main Module

We display the prompt used in the system’s Main Module. A single annotated exemplar, i.e., $f(T_1, Q_1) || P_1$, is presented in Figure 2. Multiple exemplars are included before the definition of the function to be completed. The final segment of the prompt, $f_{inc}(T, Q)$, corresponding to an example query, is shown in Figure 3.

The column description notably includes (1) the number of non-null values present in each column, (2) the type of the column as parsed by Pandas, (3) a list of all Python types present in the column, and (4) example values of categorical columns enclosed in quotes to highlight cases requiring special handling. For instance, all categorical values in the column ‘Name’ in Figure 2 begin with a space character. In all our examples, we include up to five values. If all values in a categorical column are fewer than five, we explicitly note their values.

The LLM is instructed, within the same prompt, to first predict the columns used to answer the question, their types, and the type of the answer. It then generates reasoning comments along with the code to derive the answer. The response of Llama 3.1 Instruct 405B for the prompt ending in Figure 3 is presented in Figure 4. The model correctly identifies the relevant columns and the answer type. Subsequently, it generates the code in multiple steps, clearly outlining its reasoning through comments above the corresponding statements.

A.2 Error-Fixing Module

An example prompt for the Error-Fixing Module concerning an execution error is displayed in Figure 5. The actual representation of roles in the implementation varies depending on the LLM used. The task description is followed by the function signature and the erroneous code generated by the LLM, $f(T, Q) || P$. The generation from the Main Module is then followed by the error message. The

model is subsequently tasked with rewriting the function to both fix the execution error and correctly answer the question. The completion of Llama 3.1 Instruct 405B for the prompt in Figure 5 is shown in Figure 6.

The error message conveys critical information that the model could not infer from the initial prompt. For example, although the response of the Main Module in Figure 4 follows the essential steps to answer the question, the model is unaware that the column ‘Weight Class’ contains the value ‘Open’. Given the potentially large size of tables in the dataset, crucial information necessary to answer the question may not be included in the prompt but can instead be inferred from the error message.

A.3 Prompt Size

M.M Avg. #Tokens		E.F.M Avg. #Tokens	
DB	DB Lite	DB	DB Lite
Llama Family (128K)			
19,160.9	6,518.2	2,738.4	3268.2
Claude 3.5 Sonnet (200K)			
22,764.7	7,391.7	6,573.4	4,037.8
Qwen2.5-Coder 7B (128K)			
21,816.1	7,058.2	2,700.2	3,117.7

Table 4: Average Token Count for the prompt generated in the Main Module (M.M) and Error-Fixing Module (E.F.M) for both subtasks of the test set for the different model families. The size of the context window for the different families is presented in parenthesis.

We present the average number of tokens in the prompts for both the Main and Error-Fixing Modules across the different model families in Table 4. On average, the Main Module utilizes approximately 20K for all models. Meanwhile, the Error-Fixing Module includes the header and the erroneous code produced in the prompt without additional exemplars demonstrating the task. As a result, the Error-Fixing Module requires fewer tokens, remaining below 4K tokens in most cases.

B Exploratory data analysis

Tables The DataBench dataset consists of a total of 80 tables (65 in the train and dev sets, and 15 in the test set) of varying sizes and diverse content. The distribution of table row counts across all sets in the DataBench task is shown in Figure 7. Table sizes range from fewer than 100 rows to several hundred thousand rows. In contrast, all tables in

```

# TODO: complete the following function. It should give the answer to: How many
    ↪ players have the position 'ST'?
def answer(df: pd.DataFrame):
    """
    #,Column,Non-Null Count,Dtype,Types of Elements,Values
    0,ID,14620,uint32,[<class 'int'>],
    1,Name,14620,category,[<class 'str'>],5 example values are [' L. Suarez', ' K.
    ↪ De Bruyne', ' Bruno Fernandes', ' A. Griezmann', ' M. Acuna']
    2,Preferred Foot,14620,category,[<class 'str'>],All values are ['Right', 'Left']
    3,Position,14610,category,[<class 'str'>, nan],5 example values are ['RS', 'RCM
    ↪ ', 'CAM', 'RW', 'LB']
    The first 5 rows from the dataframe:
      ID      Name Preferred Foot Position
    0  176580      L. Suarez      Right      RS
    1  192985    K. De Bruyne      Right      RCM
    2  212198  Bruno Fernandes      Right      CAM
    3  194765    A. Griezmann      Left       RW
    4  224334      M. Acuna      Left       LB
    """
    df.columns = ['ID', 'Name', 'Preferred Foot', 'Position']
    # The columns used to answer the question: ['Position']
    # The types of the columns used to answer the question: ['category']
    # The type of the answer: number
    # Create a boolean mask for rows where Position is 'ST'
    is_st_position = df['Position'] == 'ST'
    # Count the rows where the mask is True
    st_player_count = df[is_st_position].shape[0]
    # Return the count of players with 'ST' position
    return st_player_count

```

Figure 2: Single annotated exemplar used in the prompt. Task description is directly sourced from the databench_eval package provided by the organizers.

```

# TODO: complete the following function. It should give the answer to: Is our
    ↪ average employee older than 35?
def answer(df: pd.DataFrame):
    """
    #,Column,Non-Null Count,Dtype,Types of Elements,Values
    0,Age,1470,uint8,[<class 'int'>],
    1,Attrition,1470,category,[<class 'str'>],All values are ['Yes', 'No']
    2,BusinessTravel,1470,category,[<class 'str'>],All values are ['Travel_Rarely',
    ↪ 'Travel_Frequently', 'Non-Travel']
    3,DailyRate,1470,uint16,[<class 'int'>],
    4,Department,1470,category,[<class 'str'>],All values are ['Sales', 'Research &
    ↪ Development', 'Human Resources']
    The first 5 rows from the dataframe:
      Age  Attrition  BusinessTravel  DailyRate  Department
    0   41         Yes  Travel_Rarely      1102         Sales
    1   49          No  Travel_Frequently      279  Research & Development
    2   37         Yes  Travel_Rarely      1373  Research & Development
    3   33          No  Travel_Frequently      1392  Research & Development
    4   27          No  Travel_Rarely      591   Research & Development
    """
    df.columns = ['Age', 'Attrition', 'BusinessTravel', 'DailyRate', 'Department']
    # The columns used to answer the question:

```

Figure 3: The part of the prompt that includes the query (T, Q) , i.e. $f_{\text{inc}}(T, Q)$


```
['Age']
# The types of the columns used to answer the question: ['number[uint8]']
# The type of the answer: boolean
# Calculate the average age of the employees
average_age = df['Age'].mean()
# Return True if the average age is greater than 35, False otherwise
return average_age > 35
```

Figure 4: Example response from Llama 3.1 Instruct 405B as component of the Main Module.

```
-- SYSTEM --
You are an assistant tasked with helping a user fix a code error. The user has
    ↳ written a function that is supposed to answer a question about a table.
-- USER --
# Help me fix the code error of the following function by rewriting it. The function
    ↳ should return the answer to the question: Are there more than 100 lifters in
    ↳ the weight class someone that weights 82kg would compete in?
def answer(df: pd.DataFrame):
    """
        #,Column,Non-Null Count,Dtype,Types of Elements,Values
        0,Lifter Name,...
        1,Age,...
        2,Weight Class,3000,category,[<class 'str'>],5 example values are ['59 kg', '83
            ↳ kg', '105 kg', '66 kg', '74 kg']
        ...
    """
    df.columns = ...
    # The columns used to answer the question: ['Weight Class']
    # The types of the columns used to answer the question: ['category']
    # The type of the answer: boolean
    # Find the weight class for someone who weighs 82kg
    weight_class = df['Weight Class'].unique()
    weight_class = [x for x in weight_class if int(x.split()[0]) >= 82]
    weight_class = min(weight_class)
    # Count the number of lifters in that weight class
    num_lifters = df[df['Weight Class'] == weight_class].shape[0]
    # Return True if there are more than 100 lifters, False otherwise
    return num_lifters > 100
# The function outputs the following error:
# __CODE_ERROR__: invalid literal for int() with base 10: 'Open'
-- ASSISTANT --
def answer(df: pd.DataFrame):
    """
        ...
    """
    df.columns = ...
```

Figure 5: The Error-Fixing Module Prompt for Llama Models. The special tokens used by the Llama Models when roles change are replaced by the name of the role between dashes.

```
# Find the weight class for someone who weighs 82kg
weight_class = df['Weight Class'].unique()
weight_class = [x for x in weight_class if x.replace('kg', '').strip().isdigit()
    ↳ and int(x.replace('kg', '').strip()) >= 82]
weight_class = min(weight_class, key=lambda x: int(x.replace('kg', '').strip()))
# Count the number of lifters in that weight class
num_lifters = df[df['Weight Class'] == weight_class].shape[0]
# Return True if there are more than 100 lifters, False otherwise
return num_lifters > 100
```

Figure 6: Example response from Llama 3.1 Instruct 405B as component of the Error-Fixing Module.

DataBench Lite contain only 20 rows, sampled from the larger dataset.

Figure 8 presents the distribution of column counts across all dataset splits. The number of columns in DataBench ranges from 3 to 123 in the dev set and up to 35 in the test set.

Questions DataBench and DataBench Lite contain the same set of questions, with the only difference being that DataBench Lite uses sampled subsets of the original tables. Example questions, along with their corresponding answers, answer data types, and the columns used in their derivation, are presented in Table 5. All questions are answered exclusively using information present in the table.

In the annotated train and dev sets, the derivation of answers involves one, two, or three columns. Figure 9 presents an overview of the column data types utilized across the dataset. The most commonly used column data types in these sets are number and category.

The dataset defines five possible answer data types:

1. **boolean**: The only type with a predefined set of values (True or False).
2. **number**: Includes both integers and real numbers.
3. **category**: A single categorical value represented as a string.
4. **list[category]**: A list containing a fixed number of categorical values.
5. **list[number]**: A list containing a fixed number of numerical values.

Both datasets contain an equal number of questions corresponding to each possible answer data type.

Use of Textual Columns We analyzed the role of textual columns in cases where they contribute to the answer in the annotated train and dev sets, identifying 30 such instances. These cases include questions where the text data type appears among the utilized column data types included in the annotations. Textual columns are used in one of the following ways:

1. **Uniqueness Tests**: The answer depends on whether the column contains unique values.

2. **Length Conditions**: The length of the text in each cell contributes to the answer.

3. **Existence and Occurrence of Substrings**: The answer is derived based on whether a cell contains a specific substring or the number of times the substring appears within the cell.

4. **Word Count**: The number of words in each cell contributes to the answer.

5. **Index-Based Lookup**: The answer consists of one or more cells from the column, selected based on a Boolean mask.

We observe that all of these tasks can be addressed solely through programmatic statements and do not require more advanced natural language processing capabilities, such as sentiment analysis of each cell of a textual column. Consequently, we focus exclusively on program generation and do not explore more sophisticated techniques for extracting information from textual columns, such as those proposed in (Cheng et al., 2023; Gemmell and Dalton, 2023).

C Example Generations of the Main Module

Responses generated by the Main Module for the questions in Table 5 are presented in Table 6. The responses were produced using Claude 3.5 Sonnet. The model first predicts the columns it will use to answer the question, along with the expected data type of the answer. It then generates the function completion, demonstrating a strong ability to follow exemplars and decompose its reasoning into smaller steps. This structured thought process allows for efficient verification and provides some confidence in assessing the correctness of the generated response.

Larger models, such as Claude 3.5 Sonnet and Llama 3.1 Instruct 405B, exhibit a strong grasp of the underlying programming language, making them less prone to syntactic errors (Wang et al., 2025). However, they are still susceptible to semantic errors such as reasoning mistakes. For instance, in the first entry of Table 6, the model incorrectly predicts the relevant column to use, despite the column name being partially mentioned in the question. Additionally, LLMs can make more subtle reasoning errors. In the final entry, for example, the Main Module incorrectly retrieves the four highest




Figure 7: Number of rows of all tables in train & dev set and in the test set.

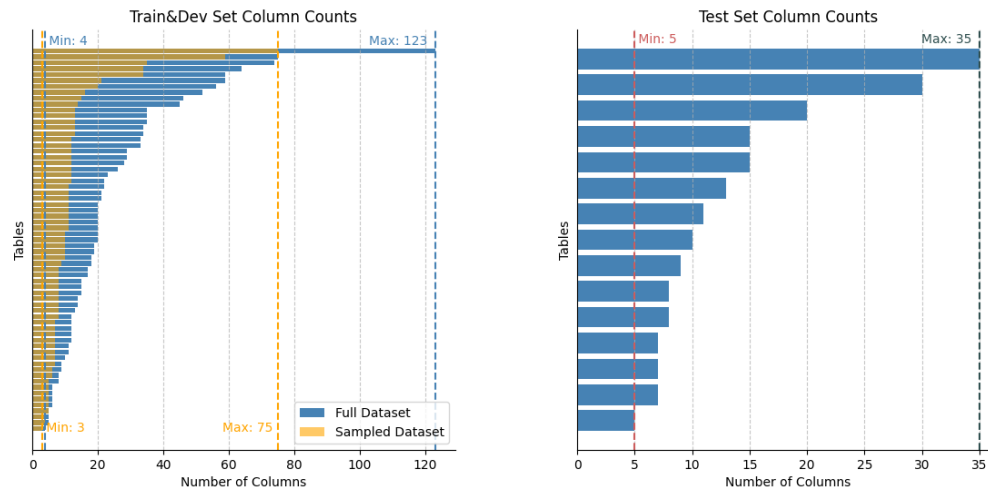


Figure 8: Number of columns of all tables in train & dev set and in the test set. Sampled versions of the tables used in the test set for DataBench Lite include all columns of the original tables.

Question	Answer	Answer Type	Columns used
How many unique types of animals are there?	4	number	['class_type']
Is the maximum level of Extraversion greater than the maximum level of Agreeableness?	True	boolean	['Extraversion', 'Agreeableness']
What are the top 3 reviewer locations with the most reviews?	['United States', 'Australia', 'Malta']	list[category]	['Reviewer_Location']
What is the most frequent age group among the respondents?	25-34	category	['How old are you?']
What are the top 4 numbers of claims in the patents?	[12, 18, 7, 13]	list[number]	['num_claims']

Table 5: Sample questions with the corresponding answers, data types of answer and the names of the columns used to answer the question sampled from DataBench Lite train and dev sets.




Figure 9: Number of times a column of each data type is used in a query. This categorization includes the samples in the train and dev set, which are annotated by the task organizers.




Figure 10: The frequencies of every answer data type in DataBench train and dev set.

values of the numerical column `num_claims` instead of identifying the four most common values, as required.

Even in cases where errors occur, the structured format of the generated responses allows for quick validation of the LLM’s output. The categorization of errors in code generation has been an area of recent research (Wang et al., 2025). Adapting these error taxonomies to the tabular setting and further refining error classification is a promising research direction. Such efforts could enhance our understanding of errors LLMs make when translating natural language queries to code in Table Question

Answering (TQA) and aid in developing methods to mitigate them.

D Model Versions

The models used in this study are the latest available versions at the time of experimentation. Specifically, they include a combination of proprietary and open-source models, along with models trained specifically for code generation. The proprietary models include Claude 3.5 Sonnet⁵, developed by Anthropic. Additionally, we incorporate open-source models such as Llama 3.3 70B⁶, Llama 3.1 8B⁷, and Llama 3.1 Instruct 405B⁸, which offer varying parameter sizes to assess performance scaling effects. Lastly, we include Qwen 2.5-Coder 7B⁹, a model specifically optimized for code generation, to evaluate its effectiveness in program synthesis tasks.

⁵anthropic.claude-3-5-sonnet-20241022-v2:0

⁶<https://ollama.com/library/llama3.3>

⁷<https://ollama.com/library/llama3.1>

⁸meta.llama3-1-405b-instruct-v1:0

⁹<https://ollama.com/library/qwen2.5-coder>

Main Module Response	Result
<pre>['animal_name'] # The types of the columns used to answer the question: [' ↳ category'] # The type of the answer: number # Count the number of unique animal names unique_animals = df['animal_name'].nunique() # Return the count of unique animals return unique_animals</pre>	Wrong Columns Used
<pre>['Extraversion', 'Agreeableness'] # The types of the columns used to answer the question: [' ↳ number[double]', 'number[double]'] # The type of the answer: boolean # Get the maximum values for both Extraversion and ↳ Agreeableness max_extraversion = df['Extraversion'].max() max_agreeableness = df['Agreeableness'].max() # Compare the maximum values and return True if ↳ max_extraversion is greater than max_agreeableness return max_extraversion > max_agreeableness</pre>	Correct Prediction
<pre>['Reviewer_Location'] # The types of the columns used to answer the question: [' ↳ category'] # The type of the answer: list[category] # Count the number of reviews per location and get the top 3 top_3_locations = df['Reviewer_Location'].value_counts(). ↳ nlargest(3).index.tolist() return top_3_locations</pre>	Correct Prediction
<pre>['How old are you?'] # The types of the columns used to answer the question: [' ↳ category'] # The type of the answer: category # Find the most frequent age group by using value_counts() most_frequent_age = df['How old are you?'].value_counts().index ↳ [0] # Return the most frequent age group return most_frequent_age</pre>	Correct Prediction
<pre>['num_claims'] # The types of the columns used to answer the question: [' ↳ number[uint8]'] # The type of the answer: list[number] # Get the top 4 values from the 'num_claims' column top_4_claims = df['num_claims'].nlargest(4) # Return the top 4 values as a list return top_4_claims.tolist()</pre>	Semantic Error

Table 6: Responses of the Main Module when queried on the questions of Table 5. Claude 3.5 Sonnet is used as the underlying LLM. The first three lines of every response correspond to the completion of the CUAT field and the remaining to the code answering the question.