

# MRT at SemEval-2025 Task 8: Maximizing Recovery from Tables with Multiple Steps

Maximiliano Hormazábal Lagos<sup>†</sup>, Álvaro Bueno Sáez<sup>†</sup>, Héctor Cerezo-Costas<sup>†</sup>,  
Pedro Alonso Doval<sup>†</sup>, Jorge Alcalde Vesteiro<sup>†</sup>

mhormazabal@gradient.org, abueno@gradient.org, hcerezo@gradient.org

palonso@gradient.org, jalcalde@gradient.org

<sup>†</sup>Fundación Centro Tecnológico de Telecomunicaciones de Galicia (GRADIANT), Vigo, Spain

## Abstract

In this paper we expose our approach to solve the *SemEval 2025 Task 8: Question-Answering over Tabular Data* challenge. Our strategy leverages Python code generation with LLMs to interact with the table and get the answer to the questions. The process is composed of multiple steps: understanding the content of the table, generating natural language instructions in the form of steps to follow in order to get the answer, translating these instructions to code, running it and handling potential errors or exceptions. These steps use open source LLMs and fine grained optimized prompts for each task (step). With this approach, we achieved a score of 70.50% for subtask 1.

## 1 Introduction

Contemporary Natural Language Processing (NLP) is limited by the volume of information (text) that can be processed effectively while maintaining contextual relevance. During response generation this constraint impacts the recall of data needed to produce correct and complete answers (Liu et al., 2024). Tabular data exemplifies this challenge in particular, since it is the day-to-day task most affected by this restriction (Ruan et al., 2024). This paper addresses the SemEval 2025 Task 8: Question-Answering over Tabular Data (Osés Grijalba et al., 2025).

In this paper we present Maximizing Recovery from Tables with Multiple Steps (MRT), a multi-step pipeline that leverages both LLMs and Python code generation to answer questions in the most factual way possible. Instead of an end-to-end strategy our system implements a sequential divide and conquer approach in which at every step either LLMs or heuristics are executed. Those steps cover from describing the tables (frequent values, column descriptions, statistical information), generating the list of instructions (in plain natural language) to carry the task and obtain the result, code execution and answer parsing.

We achieve 70.50% accuracy in the Databench Challenge test set using this approach. The code that generated these results is publicly available<sup>1</sup>.

## 2 Background

Question answering (QA) focuses on retrieving accurate answers from (Wang et al., 2025) data sets. Recent methods for QA on tabular data, such as TAPAS (Herzig et al., 2020), integrate transformers with architectures specifically tuned to extract answers directly from the tables used as context. However, LLMs have also been employed in zero-shot or few-shot strategies, since they are able to respond with a certain quality due to their prior knowledge and thus reduce the need for domain-specific fine-tuning of each domain (Kadam and Vaidya, 2020). Recent LLMs have demonstrated emergent reasoning capability, but still present difficulties with complex queries involving multiple columns, large tables, or ambiguous interpretations of a question.

Another approach is to parse natural language queries and transform them into formal queries such as SQL. Systems such as Seq2SQL (Zhong et al., 2017) or TableGPT2 (Yang et al., 2024b) are designed to generate SQL queries from relational database queries or Python code, respectively. These methods offer advantages such as greater flexibility, as they are theoretically independent of the table size (which might not fit entirely in the context window of an LLM), and greater transparency, by including an intermediate step that allows auditing and reviewing the generated queries.

To evaluate these models, reference datasets have been critical. Wikipedia-based sets, such as WikiSQL (Zhong et al., 2017) and TabFact (Chen et al., 2020), provide structured evaluation environments but do not reflect the heterogeneity of real-world tabular data (Hwang et al., 2019). In

<sup>1</sup>[https://github.com/Gradient/MRT\\_TableQA/releases/tag/v1.0.0](https://github.com/Gradient/MRT_TableQA/releases/tag/v1.0.0)

response, DataBench (Osés Grijalba et al., 2024) has been developed, which brings together 65 real-world datasets with more than 1,300 manually crafted question-answer pairs across multiple domains.

Works such as TableRAG (Chen et al., 2024) propose the use of RAG systems for tabular comprehension tasks, such as QA, employing techniques such as query expansion and a double transformation to query languages. This process translates, on the one hand, the schema to be interacted with and, on the other hand, the operation necessary to identify the cells with the answer. Also noteworthy are proposals such as Chain-of-Table (Wang et al., 2024), which implements Chain-of-Thought as an iterative reasoning mechanism. Instead of executing code in one shot, operations are executed in each iteration to add or discard information from the table until the answer is found.

Despite recent progress, some challenges still remain, such as improving reasoning within multiple rows, handling different domains and languages or integrating several tables and increasing the explainability of the full process.

### 3 System Overview

The strategy developed in our system consists of loading the table as a Pandas dataframe, and then, with the use of LLMs, generating Python code to interact with the tabular data to finally obtain the answer to each question. For this, we implemented multiple modules that are executed sequentially for each question. Some of these steps are heuristics, whereas others are LLM-based.

Each of the modules can use different LLMs. For this work, we have used multilingual models from the families Qwen2.5 (Yang et al., 2024a), Qwen2.5-coder (Hui et al., 2024), Llama-3 (AI@Meta, 2024) and Phi-4 (Abdin et al., 2024) among others.

The Figure 1 depicts an overview of the workflow of the system. The first step is understanding and analyzing the information of the table (type of data, appearance of null values, etc.). Then, an LLM generates textual instructions with the reasoning steps to follow in order to get the answer. After that, a code generation model converts the instructions to Python code. Then, the Runner executes this code. If an exception occurs during code execution or answer parsing, the system steps back

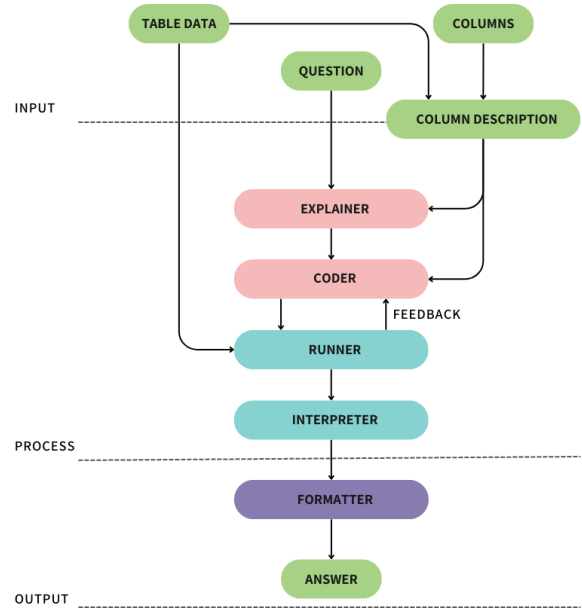


Figure 1: Diagram of the system showing all the steps involved in the generation of the response.

into the Coder in an iterative looping process until it gets a valid answer or a limit is exceeded. Finally, there are formatting steps that implement functions such as getting the answer in the desirable data type or selecting the correct number of decimals in numbers.

#### 3.1 Column Descriptor

This module aims to analyze and understand the content of the table. First, it analyzes the input table obtaining some statistical data for each column, such as the data type, the number of unique values, if it has missing values, the *max*, *min*, *mean* values and *standard deviation* (when it applies), and the most frequent values.

The second step involves serializing a subset of the table and prompting an LLM to describe the content of each column. While column names are usually descriptive, they can sometimes lack uniqueness, contain abbreviations, or be better explained within the context of the other columns.

The results of this module for each table are cached and hence this step is skipped for the following questions related to the same table.

Examples of the output of this module are shown in Listing 2 in Appendix I.

#### 3.2 Explainer

The Explainer module prompts an LLM to break down the steps required to answer a question using the table information. These instructions must be

written in natural language. The prompt includes relevant details extracted by the Column Descriptor, such as column's name, description, value type, and whether it has missing values. For numeric data types, it includes their range, and for categorical types, it lists their values if there are fewer than a configured number of unique options (fixed to 7), or otherwise just the most frequent values

The range of possible values is relevant for many questions that involve filtering by specific conditions. For example, to filter rows referring to a woman, the 'Gender' column might have various entries indistinctly like *woman*, *W*, *female*, *F*, etc. The same variety is observed in boolean values.

Guidelines are included in the prompt to force the system to use the exact given name of the columns, avoid the use of enumerations or to omit writing any code example.

The Explainer module includes a second step that prompts the LLM to review and refine the generated instructions. This process can help eliminate unnecessary steps or simplify them for greater precision.

Finally, the module parses the response to produce a list of strings, each representing an individual instruction.

Listing 3 in Appendix I contains examples of the explainer output.

### 3.3 Coder and Runner

The Coder module uses an LLM to generate Python code using the Pandas package that implements the natural language instructions in a method with the following header:

```
def parse_dataframe(df: pd.DataFrame) \
-> str:
    ...
```

The prompt includes guidelines to avoid exceptions, such as using the exact column names, casting specific data types, generating a single Python method, and avoiding the Pandas *groupby* function. The latter directive is based on empirical observations that the models we used tended to overuse this instruction, frequently resulting in numerous errors during code execution.

The LLM response is processed by a parser that employs heuristics to verify and standardize the various syntax generated by the model. As heuristics we employ different already preexisting libraries such as *autopep8*, *autoflake*, and *lib\_23* to fix minimal inconsistencies in the Python code syntax. In

particular *lib\_23* is used to parse python 2 code into python 3, and will check for missing commas/parentheses (for example). We also employ the AST tree parsing to detect when something doesn't have python code format. When detected (via parsing or exception), the system makes up to four attempts to correct them by returning the response to the LLM for revision.

Finally, the Runner module executes the code generated by the Coder and returns the result. If an exception occurs during execution, the process reverts to the Coder, with a maximum of three retries, to regenerate the code. The exception is added to the prompt to prevent it in subsequent iterations.

### 3.4 Interpreter

The Interpreter module checks if the format of the answer matches the expected type of data for the question. To achieve this, it first consults an LLM to determine the most suitable type of data to answer the question given the accepted types of the task: *Boolean*, *String*, *Number*, *List of Strings*, and *List of Numbers*.

Then, in a second call to the LLM, asks it to fit the answer to the given format if it is not already correct. With this, we correct many errors like returning numbers of booleans casted to strings (see examples in Table 7 in Appendix I).

### 3.5 Formatter

The last module in the workflow is the formatter. This module, based on rules is in charge of setting the answer in the most suitable format to match the expected task output. (see examples in Table 8 in Appendix I). For example, checks the data type of the answer, and casts it or make some format transformations.

In most cases, this module does not require any modifications due to the correction performed in previous steps. However, in certain instances, adjustments are necessary to ensure alignment with the gold labels during task evaluation.

## 4 Experimental Setup

The experimental setup consisted of two approaches for executing the modules plus the combination of configurations of each of the modules. Initially, the modules were run in series, as lighter models that could fit concurrently in memory were used. In this approach, each question for each ta-

ble was processed sequentially through all system modules. However, during testing phase with heavier models, it became necessary to implement a system that allowed to load and unload the models as it was required by each step. This approach executes each of the step in batches. All the questions are processed for each step before passing to the next one. Hence only the model used in each step is loaded in memory. This allows to load bigger models for each of the steps whilst using the same GPU without involving excessive overhead in loading/unloading models.

The hardware used to run the tests was an NVIDIA RTX-a6000 that combines 84 second-generation RT cores, 336 third-generation Tensor cores, and 10,752 CUDA cores with 48 GB of graphics memory for performance.

In 4.2, we define the model configurations used in the test phase. During development, we also used reduced versions of these models with 8B parameters.

#### 4.1 Dataset splits

Although no training of any model has been performed, the splits of the dataset are shown below.

Split	Tables	Questions
train	49	988
dev	16	320
test	15	522

Table 1: Distribution of number of tables and questions for each split in the dataset

Train, dev and test splits have been used for the development of the modules.

#### 4.2 Models

Llama 3<sup>2</sup>, Phi<sup>3</sup> and Qwen<sup>4</sup> models of different sizes have been used for the different modules of the system.

For Llama 3 only the 8B size model was used<sup>5</sup>. For Phi-4, the 14B version was chosen<sup>6</sup> and, finally, the main family of models used in the tests is Qwen. Two types of Qwen models have been executed: Qwen25 and Qwen25 code.

<sup>2</sup><https://huggingface.co/meta-llama>

<sup>3</sup><https://huggingface.co/microsoft>

<sup>4</sup><https://huggingface.co/Qwen>

<sup>5</sup><https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

<sup>6</sup><https://huggingface.co/microsoft/phi-4>

For Qwen25 two sizes have been selected: 7B<sup>7</sup> and 14B<sup>8</sup>. For Qwen25 code also the same two sizes were used: 7B<sup>9</sup> and 14B<sup>10</sup>.

Let us emphasize that the 8B models have been used mainly in the first battery of tests and development whilst the 14B models were used in the final execution of the system.

#### 4.3 Test and configuration

Table 2 summarizes the different experiments performed, with the model used in each module for each experiment.

Explainer	Coder	Interpreter
<i>llama3<sub>8B</sub></i>	<i>qwen25<sub>14B</sub>_code</i>	<i>qwen25<sub>14B</sub></i>
<i>phi4<sub>14B</sub></i>	<i>qwen25<sub>14B</sub>_code</i>	<i>qwen25<sub>14B</sub></i>
<i>qwen25<sub>14B</sub></i>	<i>qwen25<sub>14B</sub>_code</i>	<i>qwen25<sub>14B</sub></i>

Table 2: Different configurations for each step that uses LLMs

Table 2 shows that, for the most part, experiments have been performed keeping the *Qwen25 14B-coder* model in the Coder module and selecting different models for the Explainer module. The *Llama*, *Phi-4*, *Qwen2.5 14B*, and *Qwen2.5 7B* models have been tested in the Explainer, whereas *Qwen2.5 7B* was selected to be the main interpreter mainly because during the development time small tests were performed in that module with both *Llama* and *Phi-4* and the results were not remarkable. Finally, note that the Column Descriptor module is not in the table to save space. The module used for Column Descriptor was *Qwen2.5 7B* in all experiments and was run separately because the column description process is done per table and not per question.

### 5 Results

#### 5.1 Performance in Validation split

Table 3 shows the accuracy of the system using different models. The model indicates the one used in the Explainer. For the Column Descriptor and Coder all the models share the *Qwen2.5 7B* and *Qwen2.5 14B-Coder* respectively. The Explainer step is performed by the *Qwen2.5 14B*. The Ensemble is obtained by the majority voting of the three

<sup>7</sup><https://huggingface.co/Qwen/Qwen2.5-7B-Instruct>

<sup>8</sup><https://huggingface.co/Qwen/Qwen2.5-14B-Instruct>

<sup>9</sup><https://huggingface.co/Qwen/Qwen2.5-Coder-7B-Instruct>

<sup>10</sup><https://huggingface.co/Qwen/Qwen2.5-Coder-14B-Instruct>



models. Ties are resolved prioritizing *Qwen2.5* answers over *Phi-4* and *Llama* models.

Models	Run	Interpret	Format
llama3 <sub>8B</sub>	0.563	0.613	0.606
phi4 <sub>14B</sub>	0.756	0.756	0.75
qwen25 <sub>14B</sub>	0.762	0.766	0.759
Ensemble <sub>max</sub>	0.778	0.772	0.766

Table 3: Accuracy of the different strategies in the development split using the outputs of the Runner, Interpreter and Formatter step.

As can be seen in the results, the heuristics to format the final predictions sometimes introduce additional errors.

If we filter the prediction by type of response requested (Table 4) we can see that the lists of items (either number or categorical) have a much higher difficulty than the singular responses. As expected boolean responses are easy to handle by the system as only two values are possible. Nevertheless the best individual model, *Qwen2.5 14B*, obtained better results in the categorical answers.

Answer Type	llama3 8B	phi4 14B	qwen25 14B	Ensemble
Boolean	0.75	0.844	0.828	<b>0.844</b>
Number	0.627	<b>0.851</b>	0.761	0.836
Categ.	0.639	0.754	<b>0.836</b>	0.787
ListNum	0.538	0.692	0.754	<b>0.769</b>
ListCat	0.476	0.603	<b>0.619</b>	0.587
All	0.606	0.75	0.759	<b>0.766</b>

Table 4: Accuracy of the different strategies per data type of the expected answer in the validation split

After performing the same analysis in the test split (5), we can see that in general all the models experience a poorer performance in all categories but the ranking are almost the same.

Answer Type	llama3 8B	phi4 14B	qwen2.5 14b	Ensemble
Boolean	0.659	0.791	<b>0.829</b>	0.814
Number	0.417	<b>0.628</b>	0.596	0.596
Categ.	0.459	0.635	0.703	<b>0.716</b>
ListNum	0.407	0.560	0.615	<b>0.626</b>
ListCat	0.417	0.514	0.542	<b>0.556</b>
All	0.480	0.642	0.665	<b>0.667</b>

Table 5: Accuracy of the different strategies per data type of the expected answer in the test split

## 5.2 Manual Error Analysis

We performed a manual error analysis of the results flagged as an error by the official evaluator for the *Qwen2.5*. The results are summarized in Table 6.

When passing from instructions to code some of them are usually omitted or not treated properly. That was always true when the user requested to resolve ties in alphabetical order. Certain operations such as "group-by" were avoided in the prompt as the Coder was less capable of consistently generating error-free code when using them. Although our solution uses the prompt to discourage the LLM to use certain functions, another interesting option is to provide alternative implementations for them.

One clear flaw of the current design is that the system fails to filter by certain values when they do not appear in the common values of the column (e.g. when asked for *Biden* the system does not know that this value appears as *Joe Biden* in the table unless it is in the frequent values given by the Column Descriptor. This accounts for 15% of the errors. This type of errors could be mitigated by linking values asked in the query with the real ones appearing in the table during a pre-processing step (e.g. after the natural language instructions are given). Formatting issues of the response are almost 10% of the remaining errors. Handling these errors requires careful implementation of additional post-processing heuristics, and it relates greatly to how the metric to measure the performance is actually implemented (e.g. rounding issues, partial match, how lists are expected, if ordering of the elements is taking into account or not, etc.).

The *others* set accounts for all unclassifiable errors, in general due to ambiguous questions or expected answers or incorrect ground truth samples.

Description	% error
Wrong cell value filtering	14.29
Wrong Instructions	37.66
Wrong code (incl. exceptions)	14.29
Formatting (transformations)	6.49
Formatting (answer type)	3.90
Others	23.38

Table 6: Manual analysis of the errors of the *Qwen14* model in the validation split

## 6 Conclusion

This paper has addressed our proposal, MRT, in response to the challenge proposed in Semeval 2025

regarding tabular data question-answering. This technique, which introduced a multi-step pipeline leveraging both LLMs and their ability for code generation, achieved a 70.50% accuracy.

Despite the competitive results, MRT is limited due to several factors, such as formatting the output correctly and some semantic ambiguities that are not interpreted correctly (e.g., double negation in the question). Nevertheless, the largest set of errors is due to an incorrect filter of the column values (either because the value type is incorrectly detected or because the value does not appear in the same way in the question and table). Additionally, MRT encounters difficulties in addressing abstract, more subjective, and less clear questions, which can be attributed to the size of the models employed.

Future work will focus on enhancing several of the modules to eliminate all accuracy losses introduced in the between-pipeline steps and adopting a less code-driven and more linguistic approach to ambiguous questions over tabular data.

## References

- Marah Abidin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.
- AI@Meta. 2024. *Llama 3 model card*.
- Si-An Chen, Lesly Miculicich, Julian Martin Eisenschlos, Zifeng Wang, Zilong Wang, Yanfei Chen, Yasuhisa Fujii, Hsuan-Tien Lin, Chen-Yu Lee, and Tomas Pfister. 2024. *Tablerag: Million-token table understanding with language models*.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. *Tabfact: A large-scale dataset for table-based fact verification*.
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. *Tapas: Weakly supervised table parsing via pre-training*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Kai Dang, et al. 2024. Qwen2.5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Wonseok Hwang, Jinyeong Yim, Seunghyun Park, and Minjoon Seo. 2019. *A comprehensive exploration on wikisql with table-aware word contextualization*.
- Suvarna Kadam and Vinay Vaidya. 2020. Review and analysis of zero, one and few shot learning approaches. In *Intelligent Systems Design and Applications: 18th International Conference on Intelligent Systems Design and Applications (ISDA 2018) held in Vellore, India, December 6-8, 2018, Volume 1*, pages 100–112. Springer.
- Tianyang Liu, Fei Wang, and Muhao Chen. 2024. *Rethinking tabular data understanding with large language models*. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 450–482, Mexico City, Mexico. Association for Computational Linguistics.
- 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.
- 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.
- Yucheng Ruan, Xiang Lan, Jingying Ma, Yizhi Dong, Kai He, and Mengling Feng. 2024. *Language modeling on tabular data: A survey of foundations, techniques and evolution*.
- Yuxiang Wang, Jianzhong Qi, and Junhao Gan. 2025. *Accurate and regret-aware numerical problem solver for tabular question answering*.
- 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. 2024. *Chain-of-table: Evolving tables in the reasoning chain for table understanding*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2024a. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Jian Yang, Jiaxi Yang, Ke Jin, Yibo Miao, Lei Zhang, Liqun Yang, Zeyu Cui, Yichang Zhang, Binyuan

Hui, and Junyang Lin. 2024b. Evaluating and aligning codellms on human preference. *arXiv preprint arXiv:2412.05210*.

Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning.

A Appendix I: Examples of submodule outputs

Examples of the outputs of the Column descriptor (Listing 2) and Explainer (Listing 3), and examples of transformations made in the Interpreter (Table 7) and Formatter (Table 8) steps.

```
1 {
2   "name": "trip_distance",
3   "type": "float64",
4   "missing_values": 0,
5   "unique": 1259,
6   "flag_binary": false,
7   "mean": 2.0519498,
8   "std": 1.6832561884020858,
9   "freq_values": null,
10  "description": {
11    "name": "trip_distance",
12    "description": "Distance of the taxi
13 trip, typically measured in miles or
14 kilometers."
15  }
16 },
17
18 {
19   "name": "Have you ever use an online
20     dating app?",
21   "type": "category",
22   "missing_values": 0,
23   "unique": 2,
24   "flag_binary": false,
25   "mean": 0.0,
26   "std": 0.0,
27   "freq_values": [
28     "Yes",
29     "No"
30   ],
31   "description": {
32     "name": "Have you ever use an online
33       dating app?",
34     "description": "Indicates whether
35       the respondent has ever used an
36       online dating application."
37   }
38 }
```

Figure 2: Examples of outputs of the Column Descriptions for two columns.

```
1 Question: "What is the primary type of
2 the Pok\mon with the highest
3   defense
4   stat?"
5
6 Explainer Output:
7 ['Sort the rows in descending order
8 based on the "defense" column',
9 'Select the row at the top of the
10 sorted list',
11 'Access the "type1" column of
12 the selected row',
13 'Return the value in the "type1"
14 column as the answer.']
```

Figure 3: Examples of outputs of the explainer

Input	Expected	Output
"False"	Boolean	False
True	Boolean	True (unchanged)
1, 21, 14	List of numbers	[1, 21, 14]
Water, Normal	List of strings	["Water", "Normal"]
["16.0", "1.0"]	List of numbers	[16.0, 1.0]
0.2748	Number	0.2748 (unchanged)

Table 7: Examples of transformations in the Interpreter

Input	Output
2.0	2
[38.0, 23.0, 39.0]	[38, 23, 39]
(1000, 2000, 3000)	[1000, 2000, 3000]
400	400 (no changes)

Table 8: Examples of transformations in the Formatter