

# QUIS: Question-guided Insights Generation for Automated Exploratory Data Analysis

**Abhijit Manatkar**  
IBM Research, India  
abhijitmanatkar@ibm.com

**Ashlesha Akella**  
IBM Research, India  
ashlesha.akella@ibm.com

**Parthivi Gupta \***  
Indian Institute of Technology Kharagpur  
parthivig@kgpian.iitkgp.ac.in

**Krishnasuri Narayanam**  
IBM Research, India  
knaraya3@in.ibm.com

## Abstract

Discovering meaningful insights from a large dataset, known as Exploratory Data Analysis (EDA), is a challenging task that requires thorough exploration and analysis of the data. Automated Data Exploration (ADE) systems use goal-oriented methods with Large Language Models and Reinforcement Learning towards full automation. However, these methods require human involvement to anticipate goals that may limit insight extraction, while fully automated systems demand significant computational resources and retraining for new datasets. We introduce QUIS, a fully automated EDA system that operates in two stages: insight generation (ISGEN) driven by question generation (QUGEN). The QUGEN module generates questions in iterations, refining them from previous iterations to enhance coverage without human intervention or manually curated examples. The ISGEN module analyzes data to produce multiple relevant insights in response to each question, requiring no prior training and enabling QUIS to adapt to new datasets.

## 1 Introduction

Exploratory Data Analysis (EDA) is the process of discovering meaningful insights from vast amounts of data, and it is a complex task requiring careful data exploration. There are various EDA techniques to uncover insights by analyzing patterns in the data. Automated Data Exploration (ADE) systems accelerate the EDA process through automation.

ADE literature includes statistics-based (Sellam et al., 2015; Ding et al., 2019; Wang et al., 2020; Ma et al., 2021, 2023) and interactive methods (Milo and Somech, 2016, 2018b; Agarwal et al., 2023; He et al., 2024), where users explore data through natural language queries or receive suggestions for subsequent actions. Visualization-based

techniques (Vartak et al., 2015; Demiralp et al., 2017; Srinivasan et al., 2018; Wu et al., 2024) offer visual insights and allow further queries. However, these methods can become resource-intensive due to extensive user interactions. Goal-oriented ADE approaches, generate insights based on predefined objectives (Tang et al., 2017; Seleznova et al., 2020; Omidvar-Tehrani et al., 2022; Laradji et al., 2023). This approach directs the exploration using predefined objectives, such as natural language goals or statistical measures of interestingness. While this reduces user interactions, it may constrain the insights to only those aligned with the predetermined goals.

ADE using reinforcement learning is studied (Milo and Somech, 2018a; Bar El et al., 2019, 2020; Personnaz et al., 2021; Garg et al., 2023; Manatkar et al., 2024) to achieve full automation. While these systems minimize user involvement, they often demand dataset-specific training and substantial computational resources, particularly as the number of features, categorical values, or patterns increases, making the process increasingly challenging.

### 1.1 Motivation

An effective EDA system exercises statistical examination with attention to data semantics, such as analyzing trends in *date* and *sales price* or examining the impact of *weather* on *flight delay*. Systems like (Demiralp et al., 2017; Deutch et al., 2022; Ma et al., 2023; Guo et al., 2024) leverage Large Language Models (LLMs) to drive the analysis based on natural language goals. Systems which use LLMs to generate relevant questions based on natural language goals (Laradji et al., 2023), drive insight discovery based on user queries (Wang et al., 2022), and interpret analysis objectives from the user’s natural language input to specify desired outcomes (Lipman et al., 2024) have also been proposed. Guiding EDA through insightful questions

\*Work done as part of internship at IBM Research, India.

enables purposeful exploration, clarifying analysis goals, and deriving actionable insights. In contrast, such a goal-oriented approach (Laradji et al., 2023) may overlook unanticipated critical findings.

## 1.2 Our Contributions

We propose a two-stage ADE system, QUIS, that fully automates the EDA process. In the first stage, QUIS generates questions based solely on the data semantics (dataset information like name, description, column names, and column descriptions) without requiring predefined objectives. In the second stage, QUIS uses statistical analysis to produce insights corresponding to the questions from the first stage. This research contributes to the following advancements

- **Question Generation (QUGEN)** module generates questions in iterations, where questions generated in previous iterations, along with their reasoning and relevant information, serve as examples for subsequent iterations. This approach helps generate unique questions with broader coverage by providing additional context and guidance to the LLM in each iteration. Our approach eliminates the dependency on manually curated examples and predefined analysis goals.
- **Insight Generation (ISGEN)** module analyzes the data using statistical patterns and classical search techniques to generate insights in response to the questions from the QUGEN module without requiring prior training. For a given question, this module provides multiple relevant insights.

QUIS offers notable benefits, including reduced dependency on expert knowledge, enhanced efficiency in the exploration process, the ability to uncover a broader range of insights from the data, and ease of use across various datasets.

## 2 Preliminaries

Although it is challenging to precisely define the notion of an insight due to variations in users' objectives, for this work, we adopt the definition of an insight consistent with previous studies (Ding et al., 2019; Ma et al., 2023). Consider a tabular dataset  $D = \{X_1, X_2, \dots, X_n\}$  where each  $X_i$  is an attribute (column) of the dataset. An insight, denoted by  $Insight(B, M, S, P)$ , consists of the following:

1. **Perspective** - A perspective consists of a tuple  $(B, M)$ .  $B$  represents the *breakdown* attribute, and  $M$  is the *measure*, referring to a quantity of interest from the table. Typically,  $M$  is of the form  $agg(C)$  where  $agg$  (measure function) is an aggregation function, like  $count()$ ,  $mean()$ ,  $sum()$ , etc., and  $C$  (measure column) is a numerical attribute of the dataset.  $B$  is the *breakdown* dimension, a column of interest from the table, for which we want to compare different values of  $M$ . For each perspective  $(B, M)$ , we can compute a view  $view(D, B, M)$  of the dataset  $D$  by grouping on  $B$  and calculating the measure  $M$  for each group. For example, computing  $view(D, Year, mean(Performance))$  is equivalent to applying the SQL query: `SELECT Year, AVG(Performance) FROM D GROUP BY Year.`
2. **Subspace** - A subspace  $S = \bigcup_i \{(X_i, y_{ik})\}$  is a set of filters that determine a subset ( $D_S$ ) of the dataset  $D$ . Each  $X_i$  is an attribute, and each  $y_{ik}$  is a corresponding value of the column  $X_i$  of  $D$ . A tuple  $(X_i, y_{ik})$  denotes that the dataset is to be filtered for rows where  $D[X_i] = y_{ik}$ .
3. **Pattern** - The pattern  $P$  represents the type of insight observed. It belongs to a predefined set of known patterns, such as trends or outliers.

The QUIS system incorporates the following insight types as candidates for our patterns:

1. **Trend** - An increasing or decreasing trend is seen in a set of values.
2. **Outstanding Value** - The largest (or smallest) value in a set of values is significantly larger (or smaller) than all other values in the set.
3. **Attribution** - The highest value accounts for a large proportion ( $\geq 50\%$ ) of the total of all values in the set.
4. **Distribution Difference** - The distribution of values in a set changes notably from one subspace to another.

As an example, consider the insight given by

- $B = Year, M = mean(Performance)$
- $S = \{(Department, "Sales")\}$
- $P = Trend$

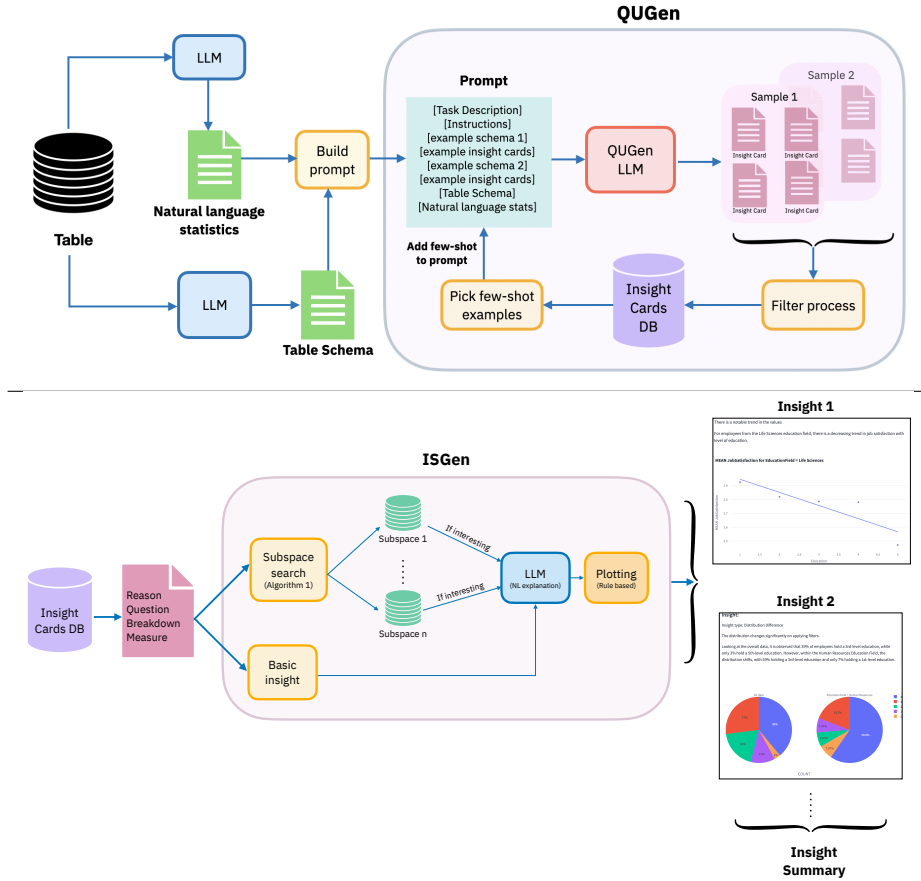


Figure 1: The Question Generation (QUGEN) module of QUIS system generates questions refined over iterations using data semantics, while the Insight Generation (ISGEN) module generates insights (bottom-right) using those questions via statistical analysis. Question is encapsulated inside the Insight Card.

This insight suggests that for the "Sales" department, there has been a trend in the average employee performance over the years.

By combining a breakdown  $B$ , a measure  $M$ , and a subspace  $S$ , we can compute a unique view of the dataset  $D$  by first applying the filters in  $S$  on  $D$  to arrive at  $D_S$ , then computing the  $view(D_S, B, M)$  as described. Let  $\mathbb{V}(D)$  be the set of all possible views of dataset  $D$  that can be computed in this manner. A search for insights involves finding views belonging to  $\mathbb{V}(D)$  for which an insight pattern  $P$  is observed. As the size of  $\mathbb{V}(D)$  grows exponentially with the number of columns in  $D$ , searching for insights by enumerating all possible views in  $\mathbb{V}(D)$  is inefficient. Therefore, it becomes important to limit the search to subspaces that are semantically meaningful and statistically relevant.

### 3 Methods

The EDA process is often guided by the questions that arise from the semantic context and the statis-

tical properties of the dataset. Hence, we propose an approach, QUIS (QUestion-guided InSight generation), that employs a two-stage process (refer to Figure 1). The first stage, QUGEN, leverages LLMs to formulate questions based on the dataset schema, basic statistics, and iteratively updated in-context examples. The second stage, question-driven insight generation (ISGEN), systematically analyzes the tabular data statistics based on the questions to uncover meaningful insights.

#### 3.1 Question Generation (QUGEN)

Our QUIS framework begins with QUGEN producing a set of Insight Cards. Each Insight Card encapsulates relevant information aligning with recent advances in automated EDA (Ding et al., 2019; Ma et al., 2021). In particular, an Insight Card (example in Figure 2) includes four components: *Question*, which is the generated natural language question aimed at guiding data analysis; *Reason*, which explains the rationale behind the generated question to help further analysis; Break-

Insight Card
<b>REASON:</b> To analyse whether there are any trends in the average performance of employees over time.
<b>QUESTION:</b> How has employee performance varied over the years?
<b>BREAKDOWN:</b> MEAN(Performance)
<b>MEASURE:</b> Year

Figure 2: Example Insight Card

down  $B$ , and Measure  $M$ . The *Reason* is used by QUGEN to enhance the coverage, and other components are used by both QUGEN and ISGEN.

QUGEN prompts the language model in a structured way to generate the Breakdown and Measure components, conditioning them on the Reason and Question. This follows the Chain-of-thought prompting approach (Wei et al., 2022), where the Reason and Question express the analysis intent behind each insight, ensuring the insights have stronger semantic justification and coherence.

### 3.1.1 Input Prompt

The prompt for QUGEN consists of several key components (for details refer to Figure 6 in Appendix), starting with a high-level description of the data analysis task objective. It then provides detailed instructions for generating an Insight Card by examining the table schema and basic statistics along with a few-shot example table schemas and their sample Insight Cards. Additionally, the prompt includes the schema of the test table and concise natural language descriptions of key statistics summarizing essential information. These statistics are generated by prompting an LLM (for prompt refer Figure 7 in Appendix) with few-shot examples to generate basic statistical questions, which are transformed into SQL, applied to the dataset, and translated into natural language responses.

### 3.1.2 QUGEN pipeline

The QUGEN LLM is prompted to generate multiple Insight Cards, as shown in Figure 1. The LLM’s response is sampled  $s$  times with a temperature  $t$ , with each sample containing  $n$  Insight Cards. However, the exact number of Insight Cards per sample may vary slightly due to the fixed output token length.

Each Insight Card undergoes a filtering process: first, cards with questions not semanti-

cally relevant to the table schema are removed using semantic similarity computed using the all-MiniLM-L6-v2 Sentence Transformers model (Reimers and Gurevych, 2019). Next, duplicate Insight Cards are eliminated based on semantic similarity between pairs of questions. Simple or rudimentary questions are filtered out by converting them to SQL queries and applying them on the dataset; if a query returns only one row, the question is discarded. This ensures that only in-depth questions are retained for comprehensive data analysis.

QUGEN is iterative in nature (refer Figure 1). It uses subset of Insight Cards generated until the current iteration as in-context examples in the prompt for the next iteration, offering supplementary context and guidance to ensure generation of unique Insight Cards distinct from that of previous iterations. A key advantage of this comprehensive approach by QUGEN module is that it eliminates the need for manually providing dataset specific in-context examples, as the Insight Cards generated by the earlier iterations help the LLM understand the dataset context during the subsequent iterations. A collection of Insight Cards accumulated over a certain number (e.g., 10) of iterations are provided as the output by QUGEN process.

## 3.2 Insight Generation (ISGEN)

This module uses classical search techniques and insight scores based on different statistical measures to identify interesting insights from the data.

To determine whether a combination of  $B$ ,  $M$ , and  $S$  reveals a particular pattern  $P$ , the module uses scoring functions based on data statistics and applies appropriate thresholds. For each insight pattern  $P$ , a corresponding scoring function  $\text{SCOREFUNC}_P : \mathbb{V}(D) \rightarrow \mathbb{R}$  is defined, along with a threshold value  $T_P$ . Further details about the scoring function and thresholds for each pattern are provided in Appendix A. If a combination of  $B$ ,  $M$ , and  $S$  results in a view  $v = \text{view}(D_S, B, M)$  such that  $\text{SCOREFUNC}_P(v) > T_P$ , the insight pattern  $P$  is considered to have been observed in  $v$ .

An Insight Card produced by QUGEN module is processed in two stages; first via identifying a basic insight followed by a subspace search for deeper insights as described below.

### 3.2.1 Basic Insight

Extraction of a basic insight helps to depict any meaningful patterns in the relationship between

$B$  and  $M$  considering the entire dataset without applying any filters. The basic insight is derived from an Insight Card by computing the view  $v_0 = \text{view}(D, B, M)$ . The applicable insight patterns are determined based on the data type of the breakdown  $B$  and the measure  $M$ . For instance, if  $B$  is an ordinal column like Year or Revenue, then the Trend pattern becomes relevant. Then, scores corresponding to these insight patterns are evaluated. For an insight pattern  $P$ , if  $\text{SCOREFUNC}_P(v_0) > T_P$ , then  $\text{Insight}(B, M, \phi, P)$  is returned as a basic insight (here  $\phi$  is an empty set).

---

**Algorithm 1** Insightful Subspace Search

---

**Require:** Dataset  $D$ , Initial subspace  $S_0$ , perspective  $(B, M)$ , language model  $LLM$ ,  $\text{SCOREFUNC}$ ,  $\text{beam\_width}$ ,  $\text{max\_depth}$ ,  $\text{exp\_factor}$

**Ensure:** Top-K subspaces by score  $\{S_1, \dots, S_k\}$

```

1: function EXPAND( $S$ )
2:    $\text{avlbl\_cols} \leftarrow D.\text{cols} - S.\text{used\_cols}$ 
    $\triangleright S.\text{used\_cols}$  are the columns used in the
   filters so far in  $S$ 
3:    $w \leftarrow \text{get\_weights}(\text{avlbl\_cols}, LLM)$ 
4:    $X \leftarrow \text{sample}(\text{avlbl\_cols}, w)$ 
5:    $y \leftarrow \text{sample}(D[X])$ 
6:   return  $S + (X, y)$ 
7: end function
8:  $\text{beam} \leftarrow [(S_0, \text{SCOREFUNC}(S_0))]$ 
9: for  $\text{depth} \in \{1, \dots, \text{max\_depth}\}$  do
10:  for  $(S, \text{score}) \in \text{beam}$  do
11:    for  $i \in \{1, \dots, \text{exp\_factor}\}$  do
12:       $S_{\text{new}} \leftarrow \text{EXPAND}(S)$ 
13:       $\text{score} \leftarrow \text{SCOREFUNC}(S_{\text{new}})$ 
14:       $\text{beam.add}((S_{\text{new}}, \text{score}))$ 
15:    end for
16:  end for
17:   $\text{beam} \leftarrow \text{top-k}(\text{beam}, k=\text{beam\_width})$ 
18: end for
19: return  $\text{beam}$ 

```

---

### 3.2.2 Subspace Search for Deeper Insights

Further insights can be generated from an Insight Card by searching for subspaces where the insight patterns are observed. To do so, we carry out a beam search procedure (Russell and Norvig, 2010) as described in Algorithm 1. The search takes an initial subspace  $S_0$ , a perspective  $(B, M)$  and a score function  $\text{SCOREFUNC}_P$  corresponding to insight pattern  $P$  as input. A beam of the current best subspaces is maintained. At each step, each sub-

space  $S$  in the beam is expanded to  $\text{exp\_factor}$  number of subspaces. Each expanded subspace  $S_{\text{new}}$  is obtained by adding a filter  $(X, y)$  to  $S$ . The selection of  $(X, y)$  happens in two steps; selecting the filter column  $X$  followed by  $y$ , the value to filter.

First, an LLM is prompted with  $(B, M)$  and an instruction to return candidate filter columns  $\mathbb{X}^{LLM} = \{X_1^{LLM} \dots X_k^{LLM}\}$  that can lead to semantically meaningful insights.  $X$  is obtained by sampling from a distribution of available columns (columns of  $D$  that have not been used in filters in  $S$ ) with the candidate filter columns  $\mathbb{X}^{LLM}$  having a probability mass of  $w_{LLM} \in [0, 1]$  distributed evenly over available columns with the rest of the mass  $(1 - w_{LLM})$  distributed over the remaining columns  $(D \setminus \mathbb{X}^{LLM})$ .  $w_{LLM}$  is decided in such a way to ensure that semantically relevant columns are picked with a high likelihood for filtering while ensuring that other columns also have a chance of being picked.

After picking  $X$ , we need to pick a value  $y$  from  $D[X]$ . To encourage the selection of values with higher frequency,  $y$  is sampled from a distribution over the unique values  $\{y_1, \dots, y_k\}$  in  $D[X]$  where the probability  $P(y_i)$  of selecting  $y_i$  is given by:

$$P(y_i) = \frac{\log(1 + N(y_i))}{\sum_i \log(1 + N(y_i))}$$

$N(y_i)$  is the frequency  $y_i$ 's appearance in  $D[X]$ .

Each candidate filter  $S_{\text{new}}$  is evaluated by calculating  $\text{SCOREFUNC}_P(\text{view}(D_{S_{\text{new}}}, B, M))$  (referred to as  $\text{SCOREFUNC}(S_{\text{new}})$  in Algorithm 1 for conciseness). After a round of expansion and evaluation, the beam is truncated to the top-k (subspace, score) pairs ranked by the score. This process repeats until the maximum desired depth of subspaces, then the final list of subspaces is returned.

The subspaces found in the search procedure are further filtered to only those  $S$  for which  $\text{SCOREFUNC}_P(\text{view}(D_S, B, M)) > T_P$  to output an insight  $\text{Insight}(B, M, S, P)$ .

### 3.2.3 Post Processing

The post-processing stage of an insight formulates the final insight response, which consists of a natural language description and a corresponding data visualization, as shown in Figure 1 (ISGEN). These components are based on the identified pattern  $P$ . For each pattern  $P$ , the natural language

response uses a predefined template to clearly communicate the key findings. For details in the plotting conditions for each pattern refer to Appendix B.

## 4 Experimental Evaluation

In our study, we evaluated the QUIS pipeline’s effectiveness using human assessment and insight scores on three datasets: Sales (Verma, 2024), Adidas Sale (Chaudhari, 2022) and Employee Attrition (Subhash, 2017). Human evaluation focused on the individual insights assessing Relevance, Comprehensibility, and Informativeness (details in Appendix C). We tested two conditions:

1. ONLYSTATS, replacing the QUGEN module with a purely statistics based card generation module, to assess the autonomous performance of ISGEN
2. QUIS, where both QUGEN and ISGEN were involved.

Replicating prior work to establish robust baselines (Ma et al., 2023; Guo et al., 2024; Weng et al., 2024) is challenging due to the lack of available code, datasets, and implementation details. Additionally, the differences in insight types and presentation formats across existing approaches make direct comparisons difficult. Therefore, our main focus is on comparing QUIS, against the baseline ONLYSTATS. For further information about the parameters of the experimental conditions, please refer to Appendix D.

The insights were evaluated by six participants who are well-versed in data analysis, with each insight assessed by three different evaluators. Each criterion - relevance, comprehensibility, and informativeness - was rated on a scale of 1 to 5; where 1 indicated the insight was not relevant, comprehensible, or informative; and 5 indicated the insight was highly relevant, comprehensible, or informative.

### 4.1 Human Evaluation

The results of the human evaluation in Figure 3 shows that for the Sales and Employee Attrition datasets, QUIS outperformed the ONLYSTATS baseline in terms of relevance, comprehensibility, and informativeness, suggesting QUIS’s overall effectiveness. However, in the Adidas Sales dataset, ONLYSTATS performed slightly better, likely due

to specific characteristics of this dataset which favour a simpler analytical approach.




Figure 3: Comparison of Average Human Evaluation Scores for QUIS and ONLYSTATS across 3 datasets.

### 4.2 Insight Score

We compare the average normalized outputs (in the range  $[0, 1]$ ) of SCOREFUNC for all insights returned by the two experimental conditions. The comparison of scores across datasets shows that QUIS consistently outperformed the ONLYSTATS condition, with higher scores across all datasets as shown in Figure 4.



Figure 4: Comparison of Insight score for QUIS and ONLYSTATS.

### 4.3 Diverse Insight Cards

To assess the effect of the iterative process of QUIS on Insight Card diversity, we analyzed the number of unique cards generated by QUIS over multiple generations (with varied number of total iterations). We started with 1 iteration and a sampling rate of 20, then progressed to 11 iterations with a sampling rate of 2, keeping the total number of outputs generated by the LLM constant at 20. In the first condition, no few-shot examples were used, while in the last condition, QUGEN iterated 10 times, appending the prompt with new few-shot examples sampled from all previous iterations (refer Figure 5).

The iterative process produced more diverse Insight Cards, as shown by the rise in the number of unique cards across successive iterations.

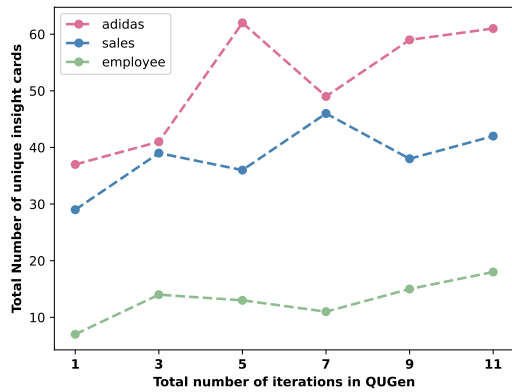


Figure 5: Total number of unique insight cards generated by QUIS under non-iterative (1 iteration) and iterative (up to 11 iterations).

## 5 Conclusion & Future Work

EDA systems often rely on user-generated, goal-oriented questions, which means the quality of the generated insights depends solely on these input questions, introducing potential overhead. To address this limitation, we propose a fully automated EDA system that generates dataset-specific questions automatically and performs insight discovery. This system operates in a data-agnostic manner, requiring no prior training, thereby minimizing the dependency on user input and streamlining the overall insight discovery process.

As a future work, we propose to enhance QUGEN to generate questions in chunks where ISGEN processes each chunk of questions before QUGEN generates the next chunk. This would enable QUGEN to use insights and their scores from previous chunks to inform the generation of subsequent chunks. Additionally, we will explore incorporating other types of insights as future work. For example, we aim to include outlier in time-series, anomaly detection, predictive insights and trend reversal to further enhance the variety and depth of insights generated by the QUIS system.

## References

Shubham Agarwal, Gromit Yeuk-Yin Chan, Shaddy Garg, Tong Yu, and Subrata Mitra. 2023. Fast Natural Language Based Data Exploration with Samples. In *Companion of the 2023 International Conference on Management of Data (SIGMOD)*, page 155–158.

AI@Meta. 2024. [Llama 3 model card](#).

Ori Bar El, Tova Milo, and Amit Somech. 2019. ATENA: An Autonomous System for Data Exploration Based on Deep Reinforcement Learning. In

*Proceedings of the 28th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 2873–2876.

Ori Bar El, Tova Milo, and Amit Somech. 2020. Automatically Generating Data Exploration Sessions Using Deep Reinforcement Learning. In *Proceedings of the 2020 International Conference on Management of Data (SIGMOD)*, pages 1527–1537.

Heemali Chaudhari. 2022. [Adidas Sales Dataset \(Adidas Sales in United States\)](#). Accessed on 19-Jul-2024.

Çağatay Demiralp, Peter J. Haas, Srinivasan Parthasarathy, and Tejaswini Pedapati. 2017. Foresight: recommending visual insights. *Proceedings of the VLDB Endowment (PVLDB)*, 10(12):1937–1940.

Daniel Deutch, Amir Gilad, Tova Milo, Amit Mualem, and Amit Somech. 2022. FEDEX: An Explainability Framework for Data Exploration Steps. *Proceedings of the VLDB Endowment (PVLDB)*, 15(13):3854–3868.

Rui Ding, Shi Han, Yong Xu, Haidong Zhang, and Dongmei Zhang. 2019. QuickInsights: Quick and Automatic Discovery of Insights from Multi-Dimensional Data. In *Proceedings of the 2019 International Conference on Management of Data (SIGMOD)*, pages 317–332.

Shaddy Garg, Subrata Mitra, Tong Yu, Yash Gadhia, and Arjun Kashettiwar. 2023. Reinforced Approximate Exploratory Data Analysis. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, volume 37, pages 7660–7669.

Yi Guo, Danqing Shi, Mingjuan Guo, Yanqiu Wu, Nan Cao, and Qing Chen. 2024. Talk2Data: A Natural Language Interface for Exploratory Visual Analysis via Question Decomposition. *ACM Transactions on Interactive Intelligent Systems (TIIS)*, 14(2):1–24.

Xinyi He, Mengyu Zhou, Xinrun Xu, Xiaojun Ma, Rui Ding, Lun Du, Yan Gao, Ran Jia, Xu Chen, Shi Han, Zejian Yuan, and Dongmei Zhang. 2024. Text2Analysis: A Benchmark of Table Question Answering with Advanced Data Analysis and Unclear Queries. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, volume 38, pages 18206–18215.

Md. Hussain and Ishtiaq Mahmud. 2019. [pyMannKendall: a python package for non parametric Mann Kendall family of trend tests](#). *Journal of Open Source Software (JOSS)*, 4(39):1556.

Maurice G. Kendall. 1975. *Rank Correlation Methods*, 4th edition. Charles Griffin, London, U.K.

William H. Kruskal and W. Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association (JASA)*, 47(260):583–621.

- Issam H. Laradji, Perouz Taslakian, Sai Rajeswar Mudumba, Valentina Zantedeschi, Alexandre Lacoste, Nicolas Chapados, David Vazquez, Christopher Pal, and Alexandre Drouin. 2023. Capture the Flag: Uncovering Data Insights with Large Language Models. In *Workshop at the Neural Information Processing Systems (NeurIPS)*.
- J. Lin. 1991. [Divergence measures based on the Shannon entropy](#). *IEEE Transactions on Information Theory*, 37:145–151.
- Tavor Lipman, Tova Milo, Amit Somech, Tomer Wolfson, and Oz Zafar. 2024. Linx: A language driven generative system for goal-oriented automated data exploration. *arXiv preprint arXiv:2406.05107*.
- Pingchuan Ma, Rui Ding, Shi Han, and Dongmei Zhang. 2021. MetaInsight: Automatic Discovery of Structured Knowledge for Exploratory Data Analysis. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD)*, pages 1262–1274.
- Pingchuan Ma, Rui Ding, Shuai Wang, Shi Han, and Dongmei Zhang. 2023. InsightPilot: An LLM-Empowered Automated Data Exploration System. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP): System Demonstrations*, pages 346–352.
- Abhijit Manatkar, Devarsh Patel, Hima Patel, and Naresh Manwani. 2024. Ilaeda: An imitation learning based approach for automatic exploratory data analysis. *arXiv preprint arXiv:2410.11276*.
- Henry B. Mann. 1945. [Nonparametric Tests Against Trend](#). *Econometrica*, 13:245–259.
- Tova Milo and Amit Somech. 2016. REACT: Context-Sensitive Recommendations for Data Analysis. In *Proceedings of the 2020 International Conference on Management of Data (SIGMOD)*, pages 2137–2140.
- Tova Milo and Amit Somech. 2018a. Deep Reinforcement-Learning Framework for Exploratory Data Analysis. In *Proceedings of the First International Workshop on Exploiting Artificial Intelligence Techniques for Data Management (aiDM)*, pages 1–4.
- Tova Milo and Amit Somech. 2018b. Next-Step Suggestions for Modern Interactive Data Analysis Platforms. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, pages 576–585.
- Behrooz Omidvar-Tehrani, Aurelien Personnaz, and Sihem Amer-Yahia. 2022. Guided text-based item exploration. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM)*, pages 3410–3420.
- Aurélien Personnaz, Sihem Amer-Yahia, Laure Berti-Équille, Maximilian Fabricius, and Srividya Subramanian. 2021. DORA THE EXPLORER: Exploring Very Large Data With Interactive Deep Reinforcement Learning. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM)*, pages 4769–4773.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3982–3992.
- Stuart J Russell and Peter Norvig. 2010. *Artificial intelligence: a modern approach*. Pearson.
- Mariia Seleznova, Behrooz Omidvar-Tehrani, Sihem Amer-Yahia, and Eric Simon. 2020. Guided exploration of user groups. *Proceedings of the VLDB Endowment (PVLDB)*, 13(9):1469–1482.
- Thibault Sellam, Emmanuel Müller, and Martin Kersten. 2015. Semi-Automated Exploration of Data Warehouses. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 1321–1330.
- Arjun Srinivasan, Steven M Drucker, Alex Endert, and John Stasko. 2018. Augmenting Visualizations with Interactive Data Facts to Facilitate Interpretation and Communication. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 25(1):672–681.
- Pavan Subhash. 2017. [IBM HR Analytics Employee Attrition & Performance \(Predict attrition of your valuable employees\)](#). Accessed on 19-Jul-2024.
- Bo Tang, Shi Han, Man Lung Yiu, Rui Ding, and Dongmei Zhang. 2017. Extracting Top-K Insights from Multi-dimensional Data. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD)*, pages 1509–1524.
- Manasi Vartak, Sajjadur Rahman, Samuel Madden, Aditya Parameswaran, and Neoklis Polyzotis. 2015. SeeDB: Efficient Data-Driven Visualization Recommendations to Support Visual Analytics. *Proceedings of the VLDB Endowment (PVLDB)*, 8(13):2182–2193.
- Shreyansh Verma. 2024. [Online Sales Dataset - Popular Marketplace Data \(Global Transactions Across Various Product Categories\)](#). Accessed on 19-Jul-2024.
- Xingbo Wang, Furui Cheng, Yong Wang, Ke Xu, Jiang Long, Hong Lu, and Huamin Qu. 2022. Interactive Data Analysis with Next-step Natural Language Query Recommendation. *arXiv preprint arXiv:2201.04868*.
- Yun Wang, Zhida Sun, Haidong Zhang, Weiwei Cui, Ke Xu, Xiaojuan Ma, and Dongmei Zhang. 2020. DataShot: Automatic Generation of Fact Sheets from Tabular Data. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 26:895–905.



Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems (NeurIPS)*, 35:24824–24837.

Luoxuan Weng, Xingbo Wang, Junyu Lu, Yingchaojie Feng, Yihan Liu, and Wei Chen. 2024. Insightlens: Discovering and exploring insights from conversational contexts in large-language-model-powered data analysis. *arXiv preprint arXiv:2404.01644*.

Yang Wu, Yao Wan, Hongyu Zhang, Yulei Sui, Wucui Wei, Wei Zhao, Guandong Xu, and Hai Jin. 2024. Automated Data Visualization from Natural Language via Large Language Models: An Exploratory Study. *Proceedings of the ACM on Management of Data (PACMMOD)*, 2(3):115:1–115:28.

## A Scoring Functions for ISGEN

Let  $\{v_1, \dots, v_k\}$  be the values for view  $v \in \mathbb{V}(D)$  for some dataset  $D$ . The following scoring functions are defined for measuring the degree to which a particular pattern is seen in  $v$ .

1. Trend - The trend pattern is observed when a sequence of values either increases or decreases monotonically. For quantifying the degree to which the trend pattern is seen, we use the Mann-Kendall Trend Test (Mann, 1945; Kendall, 1975). Specifically we use the implementation in the `pyMannKendall` package (Husain and Mahmud, 2019). Let  $MK(v)$  return the p-value calculated using the Mann-Kendall test for a  $v \in \mathbb{V}(D)$ . Then the score function is given by:

$$\text{SCOREFUNC}_{\text{Trend}}(v) = 1 - MK(v)$$

The threshold  $T_{\text{Trend}}$  is set to 0.95 so that only views having a p-value  $< 0.05$  are returned.

2. Outstanding Value - The outstanding value pattern is observed when the largest (or most negative) value is much larger (or more negative) than other values. For this pattern, the scoring function calculates the ratio between the largest value in the set and the second largest value in the set. Let  $v_{max_1}$  and  $v_{max_2}$  be the two largest (absolute) values in the set. The score is then defined as:

$$\text{SCOREFUNC}_{\text{OV}}(v) = \frac{v_{max_1}}{v_{max_2}}$$

The threshold for this pattern is set to  $T_{\text{OV}} = 1.4$

3. Attribution - The attribution pattern is observed when the top-value in a set of values accounts for more than 50% of the sum of all values. The score function used for this insight uses the ratio of the largest value to the sum of all values.

$$\text{SCOREFUNC}_{\text{Attr}}(v) = \frac{\max(\{v_1, \dots, v_k\})}{\sum_i v_i}$$

As this pattern holds when the highest value is more than 50% of the total, the threshold is set as  $T_{\text{Attr}} = 0.5$ .

4. Distribution Difference - This insight pattern can only be observed when the aggregation in the measure is `COUNT()`. Let  $v^I$  and  $v^F$  be the initial and final views. We use the Jensen-Shannon divergence (Lin, 1991) to compare the difference between the two distributions:

$$\text{SCOREFUNC}_{\text{DD}}(v^I, v^F) = JSD\left(\frac{v^I}{\sum_i v_i^I} \parallel \frac{v^F}{\sum_i v_i^F}\right)$$

The threshold is set to  $T_{\text{DD}} = 0.2$ .

## B Plotting per Pattern

- Trend: Scatter plots with trend lines are used to describe the increasing or decreasing nature of the data.
- Outstanding Value: Bar charts are used for depicting the difference in the factors.
- Attribution: Bar charts are used to show the percentage contribution of different factors
- Distribution Difference: Pie charts are used to compare the distributions before and after a condition.

## C Human Evaluation Criteria

The participants in our user study were asked to rate each generated insight on the following criteria on a scale of 1-5.

- Relevance: To what extent the insight is applicable and useful in a given context?
- Comprehensibility: To what extent is this insight understandable and easy to follow?
- Informativeness: Does the insight provide substantial information for understanding the data?

## D Experimental Conditions

### D.1 ONLYSTATS

The ONLYSTATS experimental condition replaces QUGEN with a purely statistical method for generating  $(B, M)$  pairs as follows. First, a random  $B$  is sampled from the list of all eligible columns of the table. This is followed by computing the Kruskal-Wallis test (Kruskal and Wallis, 1952) of association between breakdown  $B$  and all possible measures  $M$  in the table. The Kruskal-Wallis test is a non-parametric variance analysis test, used to determine if two sets of samples come from different distributions. The top 20 pairs of  $(B, M)$ , ranked according to the strength of association measured by the Kruskal-Wallis test are selected as input to ISGEN.

### D.2 QUIS

For QUIS, the following parameter values were used:

#### QUGEN

- LLM: Llama-3-70b-instruct (AI@Meta, 2024)
- Sampling temperature  $t = 1.1$
- Number of samples at each iteration  $s = 3$
- Number of iterations  $n = 10$
- Number of in-context examples = 6

#### ISGEN

- beam\_width = 100
- exp\_factor = 100
- max\_depth = 1
- $w_{LLM} = 0.5$


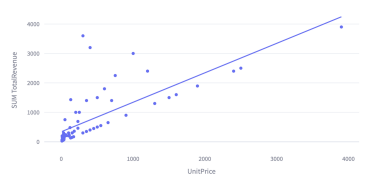
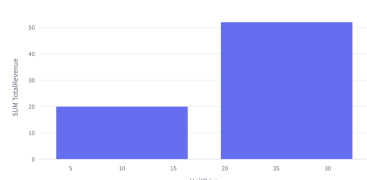
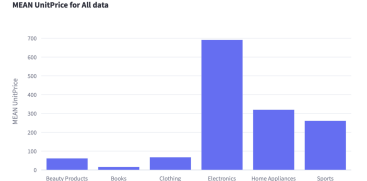
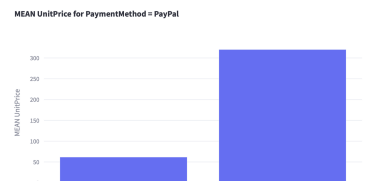
Sales Dataset		
Schema	Sample Questions	Insights
Sales ( Retailer CHAR Region CHAR SalesMethod CHAR Product CHAR PricePerUnit INT UnitsSold INT TotalSales INT OperatingProfit INT OperatingMargin DOUBLE )	Do products with higher unit prices result in higher total revenue?	<p><b>Basic Insight:</b></p> <p><b>Insight Type: Trend</b></p> <p>Total Revenue shows an increasing trend with the unit price.</p> <p><b>SUM TotalRevenue for All data</b></p>  <p><b>Insight:</b></p> <p><b>Insight Type: Trend</b></p> <p>Sales revenue paid via credit card shows an increasing trend with the increase in unit price</p> <p><b>SUM TotalRevenue for PaymentMethod = Credit Card</b></p>  <p><b>Insight:</b></p> <p><b>Insight Type: Attribution</b></p> <p>The sales of Dune by Frank Herbert at a unit price of 25.99 account for 72.23% of the total revenue generated from this product.</p> <p><b>SUM TotalRevenue for ProductName = Dune by Frank Herbert</b></p> 
	What is the average pricing strategy employed for each product category?	<p><b>Basic Insight:</b></p> <p><b>Insight Type: Outstanding Value</b></p> <p>The mean unit price is significantly higher for Electronics compared to the other product categories.</p> <p><b>MEAN UnitPrice for All data</b></p>  <p><b>Insight:</b></p> <p><b>Insight Type: Attribution</b></p> <p>The average unit price of Home Appliances is significantly higher when paid through PayPal compared to other product categories</p> <p><b>MEAN UnitPrice for PaymentMethod = PayPal</b></p> 

Table 1: QUIZ Results for Sales Dataset

Employee Attrition Dataset		
Schema	Sample Questions	Insights
Employee Attrition ( Age INT Attrition CHAR BusinessTravel CHAR DailyRate INT Department CHAR DistanceFromHome INT Education INT EducationField CHAR EmployeeCount INT EmployeeNumber INT EnvironmentSatisfaction INT Gender CHAR HourlyRate INT JobInvolvement INT JobLevel INT JobRole CHAR JobSatisfaction INT MaritalStatus CHAR MonthlyIncome INT MonthlyRate INT NumCompaniesWorked INT Over18 CHAR OverTime CHAR PercentSalaryHike INT PerformanceRating INT RelationshipSatisfaction INT StandardHours INT )	What is the relationship between employees' Education levels and their Attrition rates?	<p><b>Basic Insight:</b></p> <p><b>Insight Type: Distribution Difference</b></p> <p>For all employees, those with an education level of 3 account for 39% of all employees who have attrited from the company.</p> <p><b>Insight:</b></p> <p><b>Insight Type: Distribution Difference</b></p> <p>Employees with 3 years of education account for the highest number, with 39% of the total. However, for human resources employees, the highest number is found among those with 3 years of education, with 59% of the total. Employees with 1 year of education account for the lowest number, with 3% of the total.</p> <p><b>Insight:</b></p> <p><b>Insight Type: Distribution Difference</b></p> <p>In the Research Director role, employees with a high school education still account for the highest number, but with a slightly lower percentage of 38.00%. Conversely, employees with a college degree account for the lowest number, but with a higher percentage of 9.00%.</p>
	What is the distribution of Attrition rates across different Departments?	<p><b>Basic Insight:</b></p> <p><b>Insight Type: Outstanding Value</b></p> <p>Across all data, the Research &amp; Development department accounts for the largest proportion of employee attrition, with 65.66% of all attrition cases belonging to this department.</p> <p><b>Insight:</b></p> <p><b>Insight Type: Attribution</b></p> <p>For employees with an education background in Life Sciences, those working in the Research &amp; Development department account for 72.28% of all attrition cases.</p>

Table 2: QUIIS Results for Employee Attrition Dataset

**Task Description :** The task is to analyze a table (presented as its schema) for the purpose of Exploratory Data Analysis. Having examined the schema, you have to generate meaningful questions, and corresponding to each question a breakdown, measure and a reason. This piece of information will be further processed to generate interesting and relevant insights from the table. An insight is interesting if it helps identify one or more of the following: Meaningful relationships between variables, trends, influence of one variable over the other, anomalies or outliers.

**Instructions :**

- 1) Understand the Schema: Review the schema carefully to understand the data structure and types of columns available.
- 2) Identify Insights: Think about the different types of insights we want to uncover, such as relationships between columns, trends or anomalies. Use the provided schema and natural language stats to identify relevant and meaningful insights.
- 3) Formulate Questions: Based on the insights, formulate questions that can reveal meaningful information. Ensure that the questions are unique, relevant and not a repetition of the examples. Do not use questions related to simple data statistics (e.g., maximum length of a column).
- 4) Identify breakdown and measure dimensions for the question: Insights are obtained when a measure is compared across a breakdown dimension. The measure is a quantity of interest expressed in terms of variables of the table. It consists of
  - A measure function (aggregation) - COUNT, MEAN, MIN, MAX
  - A measure column - a numerical column of the table
 The breakdown dimension is a variable of the table across which we would like to compare values of measure to obtain meaningful insights. If the breakdown or measure dimension is absent in the question, generate relevant and related dimensions from the schema which can help provide a good insight.
- 5) Formulate a Reason: Explain what makes the question insightful and mention the reason for why the selected measure and breakdown can give a good insight. Explain why the combination of the question, breakdown and measure can help identify meaningful relationships between variables, or showcase trends, or identify outliers/anomalies.
- 6) Use [INSIGHT] Tags: Format each question using the [INSIGHT] and [/INSIGHT] tags.

**Examples :**

EXAMPLE 1:  
[EXAMPLE TABLE 1 SCHEMA]

[OUTPUT]  
Insight Card 1  
Insight Card 2  
[/OUTPUT]

EXAMPLE 2:  
.....

**Test Dataset :**

This is the information for the dataset you have to work on:

Schema  
[Test Table SCHEMA]

NATURAL LANGUAGE STATS:  
- Two payment methods  
- .....

EXAMPLE 1: .....

Please proceed to generate 10 unique and insightful questions based on the provided schema and instructions.

Figure 6: QUGen Prompt Template

**Basic Statistical Questions**

What statistical metrics would you like to know about the following database?

Example Schema (zomato):  
[STAT] What is the name of the restaurant with high number of reviews? [/STAT]  
[STAT] What is the name of the restaurant with the most diverse cuisine? [/STAT]  
[STAT] What are the different cuisines present? [/STAT]  
[STAT] What are the total number of tables in hotels and airbnbs? [/STAT]

Here is the schema to use:  
\$TABLE\_SCHEMA

INSTRUCTIONS:  
- list the stats within the [STAT] and end with [/STAT] tags, e.g:  
[STAT] How many restaurants are in the table? [/STAT]  
- Don't write anything other than the STAT

Figure 7: Natural Language Statistics Prompt Template.