

InsightEval: An Expert-Curated Benchmark for Assessing Insight Discovery in LLM-Driven Data Agents

Zhenghao Zhu^{1,*}, Yuanfeng Song^{2,*}, Xing Chen², Chengzhong Liu¹,
Yakun Cui¹, Caleb Chen Cao^{1,†}, Sirui Han^{1,†}, Yike Guo¹

¹The Hong Kong University of Science and Technology, Hong Kong, China

²ByteDance, China

Abstract

Data analysis has become an indispensable part of scientific research. To discover the latent knowledge and insights hidden within massive datasets, we need to perform deep exploratory analysis to realize their full value. With the advent of large language models (LLMs) and multi-agent systems, more and more researchers are making use of these technologies for insight discovery. However, there are few benchmarks for evaluating insight discovery capabilities. As one of the most comprehensive existing frameworks, InsightBench also suffers from many critical flaws: format inconsistencies, poorly conceived objectives, and redundant insights. These issues may significantly affect the quality of data and the evaluation of agents. To address these issues, we thoroughly investigate shortcomings in InsightBench and propose essential criteria for a high-quality insight benchmark. Regarding this, we develop a data-curation pipeline to construct a new dataset named InsightEval. We further introduce a novel metric to measure the exploratory performance of agents. Through extensive experiments on InsightEval, we highlight prevailing challenges in automated insight discovery and raise some key findings to guide future research in this promising direction.

1 Introduction

In a data-driven world, it is increasingly significant to understand and interpret vast and structured datasets (Qin et al., 2025; Lu et al., 2025). To uncover meaningful insights, data analysts require not only considering the apparent information, but also summarizing deeper patterns and relationships embedded within the dataset. Before the era of large language models (LLMs), common approaches largely relied on data processing libraries such as Pandas, NumPy, and Jupyter Notebook (Yin et al.,

2023). Therefore, insight discovery is primarily dependent on extensive manual analysis and specialized domain knowledge. Currently, LLMs have promoted the development of agent-based systems for automated data analysis and insight extraction. Recent works such as InsightPilot (Ma et al., 2023) and InsightLens (Weng et al., 2025) enable interactive data exploration via natural language, helping users rapidly identify key information.

However, comprehensive benchmarks for evaluating agents’ insight exploration capabilities remain insufficient. As one of the few publicly available datasets in this domain, InsightBench (Sahu et al., 2025) exhibits significant flaws and inconsistencies, underscoring the need for a higher-quality, more comprehensive benchmark for insight discovery. Therefore, we performed a thorough analysis of the existing Insight dataset and found numerous latent issues and inconsistencies. For example, the dataset contains substantial missing information, and some goal definitions are poorly conceived and overly broad. Several questions mention features or column names that are not in the source tables. In addition, the evaluation framework to assess agent performance is not sufficiently comprehensive, as it overlooks the accuracy of the generated insights and the novelty of the capability.

To address the issues mentioned above, we performed a deep analysis of existing deficiencies and designed three essential criteria for a high-quality insight benchmark. Moreover, we suggest two new insight types, Evaluative and Exploratory. Guided by these principles, we designed a dataset construction pipeline: (1) Refine the original goal. (2) Verify existing questions and generate new questions. (3) Answer questions and generate insights. (4) Summarize insights. Through the pipeline, we constructed InsightEval, a novel benchmark comprising 1000 insights covering six types, outperforming existing benchmarks like InsightBench. We enforced rigorous quality controls by combining

*These authors contributed equally to this work.

†Corresponding authors.

both automated checks and expert review.

In the evaluation phase, the previous assessment only relies on a single biased evaluator while ignoring erroneous outputs, and cannot recognize novel insights. Therefore, we adopted insight recall and insight precision measurement, and proposed a new metric called Insight $F1$ to comprehensively assess the agent’s ability to discover insights. Furthermore, we introduced a novel metric to evaluate the ability to uncover previously unannotated insights.

We benchmarked numerous baselines, including popular models and agent frameworks. Our results reveal that the Insight $F1$ Score can better reflect the agent’s insight discovery ability, and our InsightEval dataset provides a more comprehensive and deeper assessment.

In summary, our contributions are as follows:

- We conduct a rigorous analysis of existing insight benchmarks and define the requirements for a high-quality insight dataset.
- We introduce InsightEval, a new benchmark specifically designed to assess agents’ data analysis and insight discovery capabilities.
- We propose a comprehensive evaluation framework that combines Insight $F1$ metrics and novelty measurement.
- Our findings demonstrate that our InsightEval and evaluation framework can provide a comprehensive, in-depth, and accurate assessment of insight discovery capability.

2 Related Work

Data Analytical Agents. Several agent systems and frameworks have been proposed in the data analysis field. In data visualization, MatPlotAgent (Yang et al., 2024) combines code and multimodal LLMs, nvAgent (Ouyang et al., 2025) and MultiVis-Agent (Lu et al., 2026) generate executable programming code to visualize data. For exploratory insight, InsightPilot (Ma et al., 2023) and InsightLens (Weng et al., 2025) leverage analytic selection and multi-agent dialogue extraction to help users uncover and organize insights via natural-language interaction effectively. DAgent (Xu et al., 2025) and an LLM-based SQL-generation approach (Pérez et al., 2025) generate SQL over databases to extract information and synthesize textual reports and insights. InsightBench (Sahu et al., 2025) proposes Agent Poirot,

a multi-agent framework that iteratively generates questions, produces executable code, and derives insights. Other work includes the LangChain Pandas framework (LangChain, 2024), AutoGen (Wu et al., 2024) for customizable agent orchestration, and the ReAct (Yao et al., 2023) prompting strategies, which use reasoning and actions to enhance LLM decision-making capabilities.

Data Science Benchmarks. In Text-to-SQL research, Spider 2.0 (Lei et al., 2024) and EHRSQL (Lee et al., 2022) introduce multi-step query workflows in general and clinical contexts, including temporal and unanswerable queries. However, NL2SQL-BUGs (Liu et al., 2025), VisEval (Chen et al., 2024), and PRACTIQ (Dong et al., 2025) address semantic-error detection, visualization, conversational and ambiguous query handling, respectively. In the code-generation and tabular-analysis domains, DS-1000 (Lai et al., 2023) and JuPyT5 (Chandel et al., 2022) supply real-world programming tasks from Stack Overflow and Jupyter notebooks paired with test-execution and DSP-based evaluation. In the area of tabular data analysis, InsightBench (Sahu et al., 2025) and InfiAgent-DABench (Hu et al., 2024) focus on insight generation, covering query formulation, answer parsing, and summarization. They also assess LLM-based agent performance through structured prompting for automated evaluation.

Our dataset provides a comprehensive and rigorous benchmark for insight discovery. It has the following key characteristics: accuracy, clarity, and comprehensiveness. Experiments have clearly shown the reliability and high quality of this dataset. We sincerely hope that this dataset can promote the development of the entire field.

3 Error Analysis of Existing Insight Datasets

3.1 Insight Discovery Task Formulation

In the insight discovery task, current methods adopt a multi-agent architecture. Given tabular data and a predefined goal as inputs, the agent then generates a series of questions aligned with the goal and resolves these questions to produce answers. Subsequently, insights are discovered from the answers and are finally synthesized into a summary. There are two advanced and distinct paradigms for obtaining key information from data when solving the questions. Pérez et al. (Pérez et al., 2025) uses SQL to extract information, while Agent Poirot

Error Type	Examples	Explanation
Ambiguous Goal (E1)	Goal: Find the discrepancy and imbalance in the distribution of incidents assigned across categories. (<i>Flag-1</i>) Goal: Identify trends and underlying factors or correlations contributing to the increase in TTR. (<i>Flag-10</i>)	These Goals are overly vague and non-actionable, lacking explicit metric definitions, analysis dimensions, and temporal scope.
Undefined Data Type (E2)	Question: Do we observe any trend in the volume of incidents? Data Type: time series (<i>Flag-4</i>) Question: How does the success rate of goals across different categories compare? Data Type: comparative (<i>Flag-35</i>)	The predefined four categories are Descriptive, Diagnostic, Predictive, and Prescriptive. The example shows undefined types.
Erroneous Questions (E3)	Question: How do the distributions ... compare across departments? (<i>Flag-80</i>) Table Columns Name: category, state, closed at, opened at, closed by, number, sys updated by, location, assigned to, caller id, sys updated on, short description, priority, assignment group	The questions raised do not match the content of the table and involve non-existent column names.
Irrational Insights (E4)	Insight: Insufficient data to identify trends in resource allocation for 'Cost Reduction' goals. (<i>Flag-29</i>) Table Schema: Column: category (object), ... values: ['Employee Satisfaction', 'Cost Reduction', 'Efficiency', 'Customer Satisfaction', 'Revenue Growth']	This insight states that the data is insufficient for analysis. However, there is a corresponding "Cost Reduction" in the table schema.
Reduplicative Insights (E5)	Insight 1: Specific hardware issues related to Printer Malfunctioning are predominantly mentioned in incident descriptions. (<i>Flag-1</i>) Insight 2: Most of the hardware incidents are related to printer issues. (<i>Flag-1</i>)	These two insights both state that hardware issues mainly revolve around printer malfunctions.

Table 1: Issues and Deficiencies in InsightBench. *Flag-1*, *Flag-10* etc. means the dataset name in InsightBench.

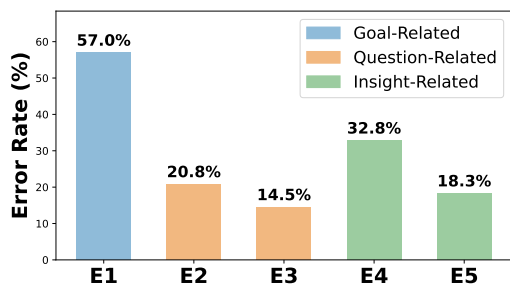


Figure 1: The proportion of each error type. E1, E2, E3, E4, E5 correspond to the Error Type in Table 1.

(Sahu et al., 2025) employs Python code and data analysis libraries to retrieve the vital information.

3.2 Preliminary Study

To establish a reliable and higher-quality benchmark, we conducted an in-depth preliminary study of the current InsightBench dataset (Sahu et al., 2025). To assess quality, we utilized GPT-4o (OpenAI, 2024) as the backbone of the multi-agent framework called Agent Poirot, randomly sampling 50 data points and generating insights. Through comprehensive analysis, we uncover numerous issues and deficiencies that undermine both data quality and the evaluation process. We list these problems in Table 1 and also calculate the proportion of each type of problem in the dataset, which is shown in Figure 1. Finally, we summarize four key observations as follows:

Observation 1: Dataset Formatting and Textual Errors. We identified multiple examples of missing or flawed content in InsightBench. Specifi-

cally, some lacked a task goal, while others were missing their associated tabular data. We consider that these omissions resulted from oversights during the dataset construction. Moreover, we encountered cases in which a question and its answer were present, but the corresponding analytical insight was omitted. Meanwhile, several questions even contained insight types that were not defined. These inconsistencies may cause deviations in the subsequent analysis and evaluation results.

Observation 2: Ambiguous and Overly Broad Goals. Among those samples, we observed that many goals were overly broad and lacked specificity. Note that a table can provide numerous potential insights through the comparative analysis of different columns. Consequently, a goal defined too generically often yields insights that are diffuse, unfocused, and lack substantive value.

Observation 3: Substandard Quality of Generated Questions and Extracted Insights. After checking the generated questions and insights, we found pervasive quality deficiencies that likely led to low evaluation scores. First, some questions refer to column names that do not exist in the table, resulting in either missing insights or logical unsoundness. Second, several insights assert insufficient information, but the requisite data were present in the table. Furthermore, we observed examples of redundant or reduplicative insights.

Observation 4: Insufficiently Comprehensive Evaluation Protocols. At present, evaluating insights is primarily based on automated text match

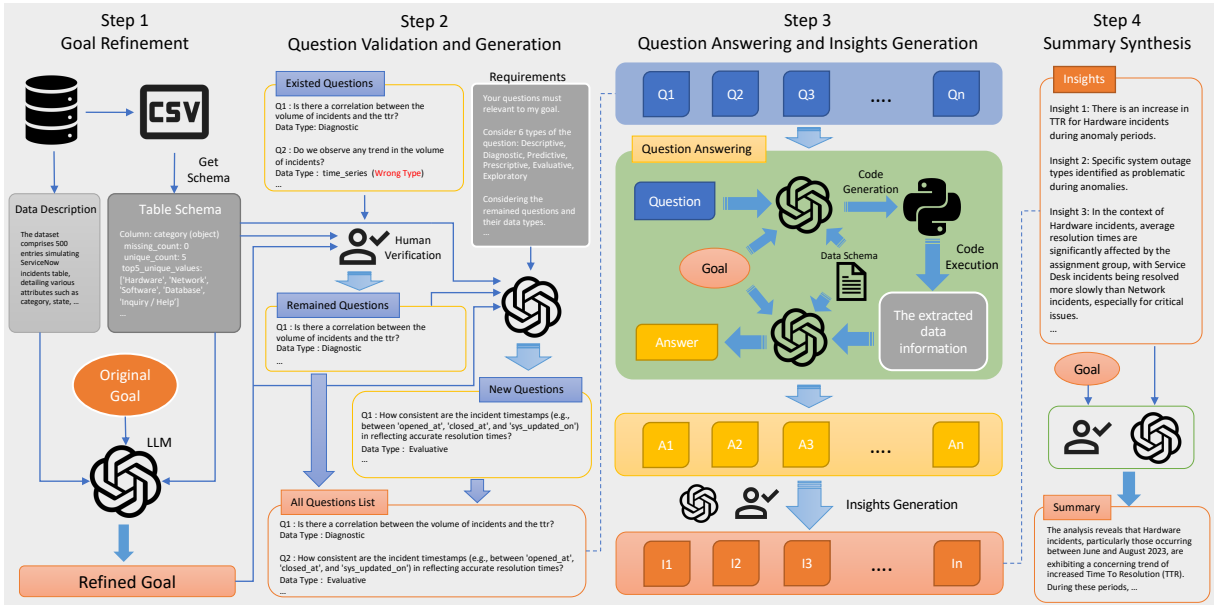


Figure 2: The dataset construction pipeline of InsightEval. This pipeline consists of 4 steps: 1) Refine the original goal through data description and schema. 2) Construct a new question list by verifying existing questions and generating new questions. 3) Extract key information by code generation and execution, then answer all the questions based on the refined goal, data schema, and extracted data, and finally generate insights. 4) Summarize all the insights by referring to the goal.

metrics and G-Eval scoring. In InsightBench, it depends exclusively on LLAMA-3-Eval as the evaluator, thereby risking the model’s inherent biases. Moreover, the evaluation merely measures how many ground-truth insights are matched by predicted insights while neglecting the generated erroneous insights. Lastly, it concentrates solely on discovering pre-annotated insights and does not recognize or reward the discovery of novel insights.

4 Dataset Construction Methodology and Evaluation Framework

4.1 Benchmark Requirements

From the observation above, we summarize three critical requirements of a High-Quality Insight Benchmark as follows:

R1. A Clearly Defined Task Goal. Each piece of data should be accompanied by a well-specified and unambiguous goal. This goal must explicitly state the intended analytical focus, including the relevant comparison metrics, dimensions of analysis, and quantitative evaluation criteria.

R2. High-Quality Questions and Insights. Each question should be tightly aligned with the tabular data and the defined goal, avoiding any speculative or unfounded formulation. Moreover, the

questions should be comprehensive, multidimensional, and reflective of diverse analytical perspectives, and the resulting insights should be meaningful, informative, and well-rounded.

R3. Multi-Perspective and Comprehensive Automatic Evaluation.

The evaluation framework should incorporate multiple LLM evaluators to mitigate individual model biases and enhance scoring reliability. In addition to comprehensively evaluating the relevance and correctness of the generated insights, the evaluation should also assess the agent’s ability to discover novel, valuable insights that are not present in the predefined ground-truth.

Based on the requirements outlined above, we constructed a new dataset for insight discovery. Each data point is accompanied by a clearly defined and unambiguous objective (**R1**), as well as meticulously curated, high-quality questions and corresponding insights (**R2**). Building upon this dataset, we introduce a comprehensive evaluation framework for insight discovery that incorporates multi-dimensional model scoring and explicitly accounts for a model’s capacity to autonomously discover novel insights (**R3**).

4.2 Data Construction

Insight Benchmark conventionally comprises four core components: a natural language objective

(*Goal*), a set of questions formulated to explore that objective (*Questions*), the insights derived from the question-answering process (*Insights*), and a summary synthesized from all insights (*Summary*). Formally, each data instance can be shown as follows:

$$x_i = (T_i, G_i, Q_i, A_i, I_i, S_i), \quad (1)$$

where T_i denotes the tabular data (with schema), G_i is the goal, Q_i is the question set, A_i and I_i are the corresponding answers and insights, and S_i is the final summary. In this paradigm, the goal and the tabular data serve as inputs, while the Insights and Summary constitute the ground-truth outputs.

Our preliminary investigation in **R1** and **R2** established that an exemplary insight benchmark should satisfy three key criteria: (1) **Clear Objectives**: Each goal must specify the exact metrics and analytical dimensions to be compared; (2) **High-Quality Questions**: Every question must be tightly related to both the Goal and the tabular data, and should be sufficiently broad and multi-perspective; (3) **High-Quality Insights**: Each insight should summarize from a rigorous analysis of the question-answer pair, yielding substantive and data-driven conclusions.

Most of the existing data science benchmarks require LLMs merely to answer isolated queries or questions (e.g., the InfiAgent-DABench (Hu et al., 2024)). Consequently, InsightBench (Sahu et al., 2025) is the closest match to our needs, offering 100 instances, 475 insights. However, we identified several structural and quality deficiencies. Accordingly, we have developed our dataset based on InsightBench, with the express aim of rectifying these issues and filling its gaps. The dataset is constructed via a hybrid pipeline combining manual inspection and LLM-assisted (o3-mini (OpenAI, 2025)) generation. An overview of this pipeline is shown in Figure 2. The detailed construction process of our dataset is described below:

Step 1: Goal Refinement. To ensure that each goal is precise and unambiguous, we implemented a meticulous validation pipeline combining LLM feedback with human review. First, given a table T_i , we extract its schema \mathcal{S}_i (e.g., column names, data types, and basic statistics). Then, an LLM refines the initial goal based on the schema:

$$\tilde{G}_i = L_{\text{goal}}(\mathcal{S}_i). \quad (2)$$

The refined goal is evaluated by human verification on three axes: alignment with the table, feasibility,

and clarity of formulation. Only goals satisfying a predefined quality threshold are retained:

$$G_i = \begin{cases} \tilde{G}_i, & H_{\text{goal}}(\tilde{G}_i, \mathcal{S}_i) \geq \theta_{\text{goal}}, \\ \text{revision,} & \text{otherwise,} \end{cases} \quad (3)$$

where $H_{\text{goal}}(\cdot)$ denotes human evaluation function.

Step 2: Question Generation and Validation.

We first manually review the existing questions and retain those aligned with the refined goal, forming an initial set $Q_i^{(0)}$. To ensure comprehensive coverage, we categorize insight questions into six types: *Descriptive*, *Diagnostic*, *Predictive*, *Prescriptive*, *Evaluative*, and *Exploratory*. For each data instance, we require exactly ten questions and enforce coverage over all insight types:

$$|Q_i| = 10, \forall \tau \in \mathcal{T}, \exists q \in Q_i \text{ s.t. } \text{type}(q) = \tau \quad (4)$$

where \mathcal{T} denotes the set of six insight categories.

Additional questions are generated by an LLM conditioned on the table schema, the refined goal, and the retained questions:

$$Q_i^{\text{new}} = L_q(\mathcal{S}_i, G_i, Q_i^{(0)}). \quad (5)$$

The final question set Q_i is obtained by selecting valid questions from $Q_i^{(0)} \cup Q_i^{\text{new}}$ and passing them through human validation.

Step 3: Answering Questions and Generating Insights.

For each question $q \in Q_i$, we generate answers and insights through a three-stage pipeline. First, LLM produces executable data analysis code:

$$C_{iq} = L_{\text{code}}(\mathcal{S}_i, G_i, q), \quad (6)$$

which is executed locally on T_i to obtain structured analytical outputs:

$$O_{iq} = \text{Exec}(C_{iq}, T_i). \quad (7)$$

Conditioned on these outputs, the LLM generates a concise factual answer:

$$a_{iq} = L_{\text{ans}}(O_{iq}, G_i, q), \quad (8)$$

followed by an interpretive insight that contextualizes the answer with respect to the goal:

$$s_{iq} = L_{\text{ins}}(G_i, \mathcal{S}_i, q, a_{iq}). \quad (9)$$

All generated answers and insights are manually verified for correctness and relevance. Invalid

or misleading outputs are removed. To eliminate redundancy, we apply a semantic de-duplication operator over the insight set:

$$I_i \leftarrow \text{Dedup}(I_i), \quad (10)$$

and repeat question and insight generation until no new unique insights are introduced, i.e.,

$$|\text{Dedup}(I_i^{(t)})| = |\text{Dedup}(I_i^{(t-1)})|. \quad (11)$$

Step 4: Summary Synthesis. Finally, given the refined goal, all validated questions, and their corresponding insights, the LLM synthesizes a comprehensive summary:

$$S_i = L_{\text{sum}}(G_i, Q_i, I_i). \quad (12)$$

The summary is subsequently verified through human review to ensure conciseness, factual consistency, and actionable value.

4.3 Evaluation Framework

Evaluating the insight discovery capabilities of the agent on InsightEval requires comparing the agent-generated insights (I) with the annotated ground-truth insights (GT). To address the shortcomings identified in **Observation 4**, we propose a set of revised evaluation criteria, as articulated in **R3**, and design novel metrics accordingly. Our automated evaluation framework employs three principal measures in insights: recall, precision, $F1$, and novelty. With these measurements, we can comprehensively assess an agent’s ability in insights discovery. Furthermore, we also perform a dedicated evaluation of summary synthesis. Below, we detail each component of our methodology.

Insights Recall Evaluation. To assess if ground-truth insights are discovered, we need to calculate the recall rate by adapting the iterative matching protocol. We count the scores between each ground-truth insight ($gt \in GT$) and each agent-generated insight ($i \in I$). Then we record the highest-scoring counterpart based on each ground-truth insight (gt) and calculate the expectation score (E) as the final output. The formula for recall evaluation is shown as in Equation 13, with S representing the evaluator.

$$\text{Score}_{\text{recall}} = E_{gt \sim \text{Unif}(GT)} \left[\max_{i \in I} \mathcal{S}(gt, i) \right] \quad (13)$$

Insights Precision Evaluation. Only focusing on the recall rate may overlook the possibility that agents generate irrelevant or unnecessary insights. To address this limitation, it is essential to further evaluate the accuracy of each generated insight to enhance the overall evaluation system. Similarly, we also enumerate the scores between the ground truth and the agent-generated insight. However, to calculate the precision rate, we need to record the highest score based on each agent-generated insight (I). The formula for precision evaluation is presented as in Equation 14.

$$\text{Score}_{\text{precision}} = E_{i \sim \text{Unif}(I)} \left[\max_{gt \in GT} \mathcal{S}(i, gt) \right] \quad (14)$$

Insights $F1$ Evaluation. To comprehensively evaluate the capability of insight discovery, we proposed a new measurement called Insight $F1$ Score. With the insight recall score and the insight precision score, we can calculate the insight $F1$ score through the formula in equation 15.

$$\text{Score}_{F1} = \frac{2 * \text{Score}_{\text{recall}} * \text{Score}_{\text{precision}}}{\text{Score}_{\text{recall}} + \text{Score}_{\text{precision}}} \quad (15)$$

Insights Novelty Evaluation. Given the limitations of merely aligning with ground-truth insights, it is essential to evaluate the capacity of discovering novel insights. We identify insights with a G-Eval score exceeding 5 in the insight precision evaluation as correct, while the other insights are classified as incorrect and subjected to a secondary evaluation focused on innovation. During the evaluation, we utilize three distinct LLMs to mitigate bias. The insight can be labeled as a potential novel insight when at least two models judge it as correct. To get more accurate judgments, we provide LLMs with contextual information, including the goal, table schema, and historical insights, and use a Chain-of-Thought (CoT) reasoning framework. The formula for novelty evaluation is expressed as in Equation 16, where $\text{LLM}_j(i) \in \{0, 1\}$, $\delta \in \{0, 1\}$, j is the number of LLMs, $\mathbf{1}$ means indicator function, M and N indicate the number of correct and incorrect insights in precision evaluation, respectively.

$$\text{Score}_{\text{novelty}} = \frac{M + \delta * C}{N + M} \quad (16)$$

where C is the count of novel samples:

$$C = \sum_{i=1}^N \mathbf{1} \left(\sum_{j=1}^3 \text{LLM}_j(i) \geq 2 \right) \quad (17)$$

Dataset	Input	Output	Data Size	Construction Method
Spider 2.0 (Lei et al., 2024)	Question	SQL Query	632 Problems	Machine & Human-Labeled
MatPlotBench (Yang et al., 2024)	Question+Table	Visual Image	100 Cases	Machine & Human-Labeled
InfiAgent-DABench (Hu et al., 2024)	Question+Table	Answer	603 Cases	Machine-Labeled
MedAgentsBench (Tang et al., 2025)	Question	Answer	862 Problems	Existed Dataset Combined
InsightBench (Sahu et al., 2025)	Goal+Table	Insights	100 Cases (475 Insights)	Human-Labeled
InsightEval	Goal+Table	Insights	100 Cases (1000 Insights)	Machine & Human-Refined

Table 2: Comparison of InsightEval with other existing benchmarks.

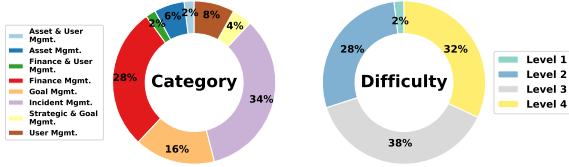


Figure 3: Data Statistics in InsightEval.

When $\delta = 1$, the formula will calculate the innovation score. For comparison, we set $\delta = 0$ to obtain the original score during the evaluation.

Summary Evaluation. For summaries, we perform a one-to-one comparison between each ground-truth summary and its generated counterpart. Then we use evaluators to score each pair to derive an evaluation of summary quality.

5 InsightEval: Statistic and Quality Analysis

5.1 Benchmark Statistic

InsightEval comprises 100 instances, each with its corresponding CSV table. For each instance, we provide 10 individual insights and one overall summary. We adopt the difficulty and category established by InsightBench. Each data point is assigned one of four difficulty levels and also annotated with six commercial analytics scenarios, with distribution shown in Figure 3. In addition, we counted the number of each insight category and calculated the average tokens for the corresponding questions and insights, which are presented in Figure 4. Moreover, compared to other well-regarded datasets, our InsightEval stands out for its large-scale and comprehensive insights, which are displayed in Table 2.

5.2 Data Quality Analysis

To further ensure data quality, we conducted an in-depth annotation combining evaluations by LLM (o3-mini (OpenAI, 2025)) and domain experts across three dimensions:

- **Correctness:** whether each question set strictly corresponds to the stated objective and the source

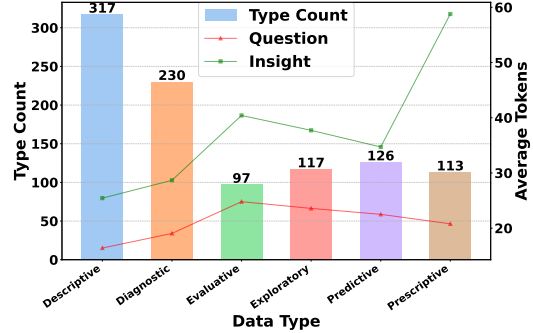


Figure 4: Data Type Distribution and Token Counts

table data without factual errors.

- **Rationality:** whether each insight satisfies the goal’s requirements and is logically sound.
- **Coherence:** whether insights are internally consistent and mutually compatible, including the overall summary’s logical flow.

We annotated all instances by both LLM and human experts, computing the accuracy rate for each dimension. The results are reported in Table 3.

Method	Correctness	Rationality	Coherence
LLM Evaluation	0.905	0.890	0.90
Human Annotation	0.935	0.920	0.95

Table 3: Results of data quality assessment.

Type	TC Similarity	Self-BLEU	Distinct-2
Questions	0.0583	0.1484	0.8677
Insights	0.0528	0.0897	0.9157

Table 4: Comparison of average TF-IDF cosine (TC) similarity, Self-BLEU, and Distinct-2 diversity scores between Question and Insight in InsightEval.

In addition, we checked each question and insight for redundancy in three metrics. First, we calculated cosine similarity over TF-IDF vector representations and averaged the resulting similarity scores. Second, we computed Self-BLEU of each sentence in questions and insights. Third, we measured Distinct-2, defined as the ratio of unique bi-grams to total bi-grams across all sentences. For TF-IDF cosine similarity and Self-BLEU scores,

Baselines	Insights Recall		Insights Precision		Insights F_1		Summary	
	ROUGE-1	G-Eval	ROUGE-1	G-Eval	ROUGE-1	G-Eval	ROUGE-1	G-Eval
<i>LLM-only</i>								
GPT-4o	0.2304	0.3389	0.2445	0.3506	0.2372	0.3447	0.2423	0.3282
DeepSeek-V3	0.2183	0.3402	0.2295	0.3554	0.2238	0.3476	0.2405	0.3332
Claude-3.7-Sonnet	0.2219	0.3265	0.2364	0.3492	0.2289	0.3375	0.2439	0.3250
<i>Single-Agent</i>								
ReAct (GPT-4o)	0.2506	0.3977	0.2573	0.4069	0.2539	0.4022	0.2654	0.3913
CodeGen (GPT-4o)	0.2488	0.4289	0.2579	0.4412	0.2533	0.4350	0.2598	0.3991
<i>Multi-Agents</i>								
DeepResearchAgent (GPT-4o)	0.2993	0.5017	0.3079	0.5198	0.3035	0.5106	0.3363	0.4279
Pandas Agent (GPT-4o)	0.3024	0.4973	0.3112	0.5133	0.3067	0.5052	0.3289	0.4021
Agent Poirot (GPT-4o)	0.2907	0.5293	0.2945	0.5487	0.2926	0.5388	0.3496	0.4334
Agent Poirot (Deepseek-V3)	0.2590	0.4984	0.2658	0.5453	0.2624	0.5208	0.3165	0.4772
Agent Poirot (Claude-3.7-Sonnet)	0.2623	0.5519	0.2673	0.6261	0.2648	0.5867	0.3178	0.4746
Avg.	0.2584	0.4411	0.2672	0.4657	0.2627	0.4529	0.2901	0.3992

Table 5: Comparative performance of various baselines at the Insights and Summary Level in InsightEval.

higher values indicate redundancy and more repetition. However, the Distinct-2 value that is closer to 1 reflects greater lexical diversity and lower redundancy. We present the redundancy statistics in Table 4. Through this process, our dataset attains a high standard of reliability and scholarly validity.

6 Experiments

6.1 Experimental Setup

Baselines. We evaluate various baselines on InsightEval using three prominent large language models currently in widespread use, such as GPT-4o (OpenAI, 2024), Deepseek-V3 (DeepSeek-AI et al., 2025), and Claude-3.7-Sonnet (Anthropic, 2025). These LLMs are employed as a backbone across multiple agent-based frameworks, including ReAct (Yao et al., 2023), CodeGen (Majumder et al., 2025), DeepResearchAgent (Zhang et al., 2025), Pandas Agent (LangChain, 2024), and Agent Poirot (Sahu et al., 2025). Details of each agentic baseline are described in Appendix D.2.

Implementation Details. We configure each agent with a temperature of 0 to ensure determinism. In Agent Poirot, we run a total of 4 rounds, with 3 new questions generated in each round. Similarly, we also have another agent generate the same number of questions to ensure the rationality.

Evaluation Metrics. For insight recall, precision, and summary assessment, we employ two evaluators: the ROUGE-1 (Lin, 2004) and G-Eval (Liu et al., 2023). Specifically, the G-Eval score is the average score across GPT-3.5-Turbo (Ye et al., 2023) and Gemini 2.5 Pro (Comanici et al., 2025).

Next, we use formula 13 and 14 to calculate the recall and precision scores. The final score in G-Eval should be normalized to facilitate a comparison with ROUGE-1. Then we calculate the insight F_1 using the formula 15. For comparison, we sampled 30 data points and scored them by ten human experts. To measure insight novelty, we utilize the formula 16 to calculate the original novelty scores and new novelty scores.

6.2 Experimental Results and Findings

Insight F_1 Score Provides a Better Reflection of Insight Capabilities. For insight F_1 results shown in Table 5, Pandas Agent achieved higher ROUGE-1 scores, whereas Agent Poirot based on the Claude 3.7 Sonnet substantially outperformed others in G-Eval.

To further analyze the results, we conducted manual expert scoring, comparing it with Insight Recall (InsightBench Metric) and F_1 , which are shown in Figure 5. Notably, Insights F_1 Scores exceeded Insight Recall Scores across all agents, and were closer to the Human Evaluation Scores. This result underscores that Insights F_1 Scores more effectively evaluate both the agent’s insight discovery capability and its alignment with human judgment.

InsightEval Provides a Comprehensive Evaluation of the Agent’s Performance. Our results reveal a range of challenges and key findings in insight discovery, as follows:

Finding 1: Agents Exhibit Limited Breadth in Insight Exploration. In Table 5, the Precision Scores of the insights outperform the recall scores, indicating the tendency that agents generate the most confidently correct insights while avoiding

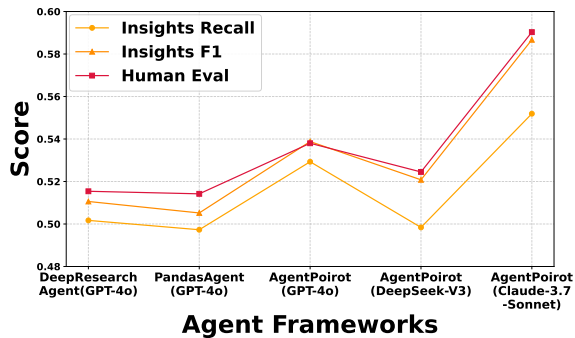


Figure 5: Comparison of G-Eval (Insight Recall & Insight $F1$) and expert human-evaluation scores.

uncertain or exploratory output. This reduces random or spurious content but results in substantial redundancy, which means highly scored insights may be correct and duplicated. Consequently, although agents demonstrate strong output quality, they come short in comprehensive exploration.

Finding 2: Multi-Agent Systems Outperform Single-Model Baselines. As shown by the baseline comparisons in Table 5, multi-agent approaches achieve superior results on both the Insight and Summary evaluations. We attribute this advantage to the ability of multi-agent design to decompose the task pipeline and perform multiple rounds of targeted analysis and exploration, which enables more thorough reasoning and information synthesis. Concretely, according to G-Eval scores, Agent Poirot (Claude-3.7-Sonnet) achieves the highest Insight $F1$, while Agent Poirot (Deepseek-V3) achieves the highest Summary score. Moreover, the relative gain is slightly larger for Insight $F1$ than for the Summary metric.

Finding 3: The Novelty of Agents Linked to the Capabilities of Backbone Models. We compared Original Novelty Scores with New Novelty Scores, as illustrated in Figure 6. We observed that all agents achieved measurable improvements in novelty. In particular, the agent built on the Claude-3.7-Sonnet achieved the highest New Novelty Score of 76.2%. We attribute this performance to its superior code-generation abilities, which likely enable it to answer questions more effectively and derive deeper insights. By calculating the improvement ratio of the new novelty score over the original novelty score, we discovered Deepseek-V3-based agent achieves the largest rate among all agents, with 13.3%. This finding suggests that agents with lower precision may compensate by producing more creative outputs, thereby

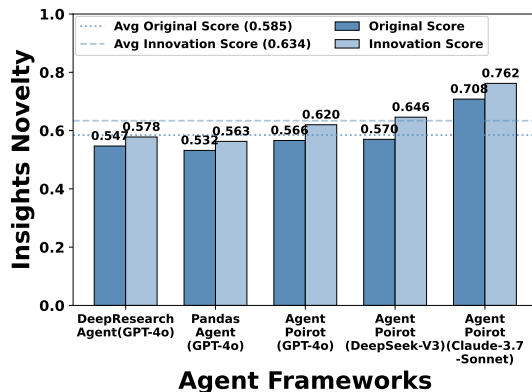


Figure 6: Performance of Novelty Evaluation.

attaining higher novelty scores.

7 Conclusion and Future Work

We present InsightEval, a novel benchmark for rigorously assessing agent’s insight discovery capabilities. Our contributions include proposing a comprehensive and high-quality dataset, developing an automated evaluation framework that spans recall, precision, and novelty, and conducting a systematic evaluation of state-of-the-art agent frameworks. Our findings uncover key challenges in automated insight discovery and offer valuable guidance for future research. In future work, we will design and evaluate state-of-the-art multi-agent systems to significantly enhance insight discovery, enabling paradigm-shifting advances in the domain.

Limitations

This work introduces InsightEval as an expert-curated benchmark for table-driven insight generation, defining a structured taxonomy, evaluation criteria, and a standardized assessment protocol validated across multiple agents and language models. However, several inherent limitations remain difficult to eliminate. i) The dataset scale and domain coverage are constrained by annotation cost and design choices, and therefore may not fully reflect the diversity of real-world settings. To address this limitation, future work will explore controlled data augmentation and synthesis to enrich under-represented scenarios, employ active and diversity-aware sampling strategies to prioritize informative and low-coverage cases, and expand domain coverage through collaborations with domain experts and community-driven data contributions. ii) The nature of an insight is intrinsically subjective, as judgments of value, actionability, and novelty vary

across users and contexts. The pure table-input setting also imposes a natural information ceiling, since many meaningful insights require external or contextual knowledge not present in the data. Ground-truth annotations are necessarily incomplete and should be viewed as representative rather than exhaustive, and novelty assessments remain time- and context-dependent as reference knowledge evolves. As a result, findings should be interpreted as relative and conditional rather than absolute. In future work, we aim to expand evaluation contexts and longitudinal settings to better understand how model capabilities evolve and adapt across domains and time.

References

- AI Anthropic. 2025. [Claude 3.7 sonnet and claude code](#).
- Shubham Chandel, Colin B Clement, Guillermo Serrato, and Neel Sundaresan. 2022. Training and evaluating a jupyter notebook data science assistant. *arXiv preprint arXiv:2201.12901*.
- Nan Chen, Yuge Zhang, Jiahang Xu, Kan Ren, and Yuqing Yang. 2024. Viseval: A benchmark for data visualization in the era of large language models. *IEEE Transactions on Visualization and Computer Graphics*.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, and 1 others. 2025. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Mingwen Dong, Nischal Ashok Kumar, Yiqun Hu, Anuj Chauhan, Chung-Wei Hang, Shuaichen Chang, Lin Pan, Wuwei Lan, Henghui Zhu, Jiarong Jiang, and 1 others. 2025. Practiq: A practical conversational text-to-sql dataset with ambiguous and unanswerable queries. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 255–273.
- Xueyu Hu, Ziyu Zhao, Shuang Wei, Ziwei Chai, Qianli Ma, Guoyin Wang, Xuwu Wang, Jing Su, Jingjing Xu, Ming Zhu, and 1 others. 2024. Infiagent-dabench: evaluating agents on data analysis tasks. In *Proceedings of the 41st International Conference on Machine Learning*, pages 19544–19572.
- Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Wen-tau Yih, Daniel Fried, Sida Wang, and Tao Yu. 2023. Ds-1000: A natural and reliable benchmark for data science code generation. In *International Conference on Machine Learning*, pages 18319–18345. PMLR.
- LangChain. 2024. [Pandas dataframe](#).
- Gyubok Lee, Hyeonji Hwang, Seongsu Bae, Yeonsu Kwon, Woncheol Shin, Seongjun Yang, Minjoon Seo, Jong-Yeup Kim, and Edward Choi. 2022. Ehrsql: A practical text-to-sql benchmark for electronic health records. *Advances in Neural Information Processing Systems*, 35:15589–15601.
- Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin SU, ZHAOQING SUO, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, and 1 others. 2024. Spider 2.0: Evaluating language models on real-world enterprise text-to-sql workflows. In *The Thirteenth International Conference on Learning Representations*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Xinyu Liu, Shuyu Shen, Boyan Li, Nan Tang, and Yuyu Luo. 2025. Nl2sql-bugs: A benchmark for detecting semantic errors in nl2sql translation. In *Proceedings of the 31th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Jinwei Lu, Yuanfeng Song, Zhiqian Qin, Haodi Zhang, Chen Zhang, and Raymond Chi-Wing Wong. 2025. Bridging the gap: Enabling natural language queries for nosql databases through text-to-nosql translation. *arXiv preprint arXiv:2502.11201*.
- Jinwei Lu, Yuanfeng Song, Chen Zhang, and Raymond Chi-Wing Wong. 2026. [Multivis-agent: A multi-agent framework with logic rules for reliable and comprehensive cross-modal data visualization](#). *Proc. ACM Manag. Data*, 4(1).
- 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: System Demonstrations*, pages 346–352, Singapore. Association for Computational Linguistics.
- Bodhisattwa Prasad Majumder, Harshit Surana, Dhruv Agarwal, Bhavana Dalvi Mishra, Abhijeetsingh Meena, Aryan Prakhar, Tirth Vora, Tushar Khot,

- Ashish Sabharwal, and Peter Clark. 2025. [Discoverybench: Towards data-driven discovery with large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- OpenAI. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- OpenAI. 2025. [Openai o3-mini system card](#).
- Geliang Ouyang, Jingyao Chen, Zhihe Nie, Yi Gui, Yao Wan, Hongyu Zhang, and Dongping Chen. 2025. [nvAgent: Automated data visualization from natural language via collaborative agent workflow](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 19534–19567.
- Alberto Sánchez Pérez, Alaa Boukhary, Paolo Papotti, Luis Castejón Lozano, and Adam Elwood. 2025. An llm-based approach for insight generation in data analysis. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 562–582.
- Zhiqian Qin, Yuanfeng Song, Jinwei Lu, Yuanwei Song, Shuaimin Li, and Chen Jason Zhang. 2025. Multi-tend: A multilingual benchmark for natural language to nosql query translation. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24632–24657.
- Gaurav Sahu, Abhay Puri, Juan A Rodriguez, Amirhossein Abaskohi, Mohammad Chegini, Alexandre Drouin, Perouz Taslakian, Valentina Zantedeschi, Alexandre Lacoste, David Vazquez, and 1 others. 2025. [Insightbench: Evaluating business analytics agents through multi-step insight generation](#). In *The Thirteenth International Conference on Learning Representations*.
- Xiangru Tang, Daniel Shao, Jiwoong Sohn, Jiapeng Chen, Jiayi Zhang, Jinyu Xiang, Fang Wu, Yilun Zhao, Chenglin Wu, Wenqi Shi, and 1 others. 2025. [Medagentsbench: Benchmarking thinking models and agent frameworks for complex medical reasoning](#). *arXiv preprint arXiv:2503.07459*.
- Luoxuan Weng, Xingbo Wang, Junyu Lu, Yingchaojie Feng, Yihan Liu, Haozhe Feng, Danqing Huang, and Wei Chen. 2025. [Insightlens: Augmenting llm-powered data analysis with interactive insight management and navigation](#). *IEEE Transactions on Visualization and Computer Graphics*.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, and 1 others. 2024. [Autogen: Enabling next-gen llm applications via multi-agent conversations](#). In *First Conference on Language Modeling*.
- Wenyi Xu, Yuren Mao, Xiaolu Zhang, Chao Zhang, Xuemei Dong, Mengfei Zhang, and Yunjun Gao. 2025. [Dagent: A relational database-driven data analysis report generation agent](#). *arXiv preprint arXiv:2503.13269*.
- Zhiyu Yang, Zihan Zhou, Shuo Wang, Xin Cong, Xu Han, Yukun Yan, Zhenghao Liu, Zhixing Tan, Pengyuan Liu, Dong Yu, Zhiyuan Liu, Xiaodong Shi, and Maosong Sun. 2024. [MatPlotAgent: Method and evaluation for LLM-based agentic scientific data visualization](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11789–11804, Bangkok, Thailand. Association for Computational Linguistics.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *International Conference on Learning Representations (ICLR)*.
- Junjie Ye, Xuanting Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, Zeyang Zhou, Chao Gong, Yang Shen, and 1 others. 2023. [A comprehensive capability analysis of gpt-3 and gpt-3.5 series models](#). *arXiv preprint arXiv:2303.10420*.
- Pengcheng Yin, Wen-Ding Li, Kefan Xiao, Abhishek Rao, Yeming Wen, Kensen Shi, Joshua Howland, Paige Bailey, Michele Catasta, Henryk Michalewski, Oleksandr Polozov, and Charles Sutton. 2023. [Natural language to code generation in interactive data science notebooks](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 126–173, Toronto, Canada. Association for Computational Linguistics.
- Wentao Zhang, Liang Zeng, Yuzhen Xiao, Yongcong Li, Ce Cui, Yilei Zhao, Rui Hu, Yang Liu, Yahui Zhou, and Bo An. 2025. [Agentorchestra: Orchestrating hierarchical multi-agent intelligence with the tool-environment-agent\(tea\) protocol](#). *Preprint*, arXiv:2506.12508.

A Details of InsightEval Benchmark

A.1 Category and Difficulty

Referring to the categories and difficulty settings in InsightBench (Sahu et al., 2025), InsightEval contains 8 business analytics categories and 4 difficulty levels. We describe the details of each category as follows:

- **Incident Management:** Tracks and analyzes operational or safety incidents to enable rapid response and root-cause investigation.
- **Asset Management:** Manages IT hardware lifecycle—procurement, deployment, maintenance—ensuring inventory visibility and optimal utilization.
- **User Management:** Maintains user profiles, roles, departments, and login status to enforce access control and support audit trails.
- **Finance Management:** Audits expense records to reveal spending patterns, optimize budget allocation, and drive cost-saving decisions.
- **Goal Management:** Monitors departmental or project objectives—planning, progress, and completion rate—to assess performance and alignment.
- **Asset & User Management:** Correlates hardware assignments with user data to optimize resource distribution and usage efficiency.
- **Finance & User Management:** Links expense transactions with individual or team activity to uncover cost behaviors and usage trends.
- **Strategy & Goal Management:** Integrates strategic plans with goal-tracking data to evaluate execution effectiveness and organizational alignment.

A.2 Insight Types

In this paper, we aim to provide a more comprehensive interpretation of data insights by extending the four insight categories originally defined in InsightBench (Sahu et al., 2025) with two additional types. A detailed description of each insight category is provided below:

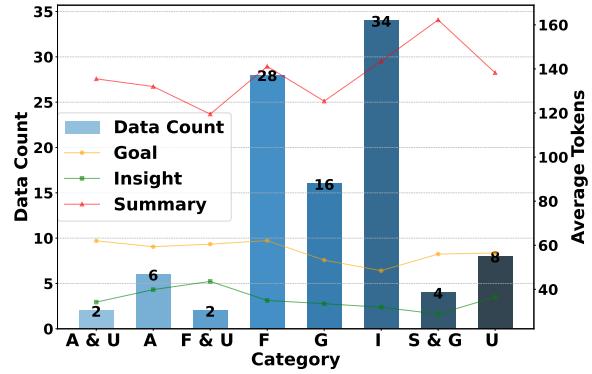


Figure 7: Category distributions with token count in different data items. A, U, F, G, I, S separately stand for Asset, User, Finance, Goal, Incident, Strategy.

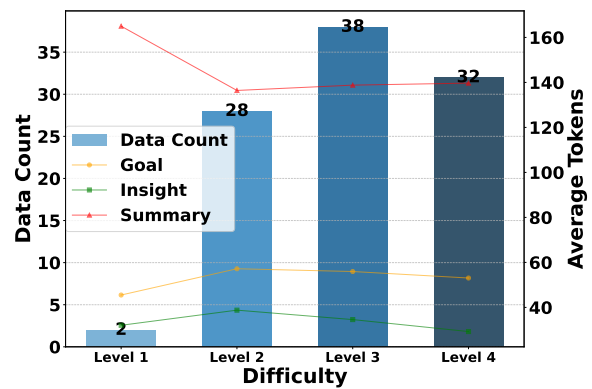


Figure 8: Difficulty distributions with token count in different data items.

- **Descriptive:** Summarize what happened. This type of analysis describes past situations by aggregating and visualizing historical data (for example, generating a chart of monthly investment portfolio returns).
- **Diagnostic:** Explain why it happened. Identifies correlations, patterns, and root causes to explain observed trends or results (for example, segmenting losses by asset category to identify key drivers).
- **Predictive:** Forecast what is likely to happen. Based on historical trends, it uses statistical models to predict future outcomes (for example, estimating the risk of default in the next quarter using credit scores).
- **Prescriptive:** Recommend specific actions to take. It suggests actionable strategies, such as optimization or risk mitigation, to achieve desired objectives (for example, advising portfolio adjustments to reduce predicted volatility).

- **Evaluative:** Assess the quality and reliability of the data and analysis. This involves evaluating the completeness, accuracy, and robustness of data and analytical methods (for example, verifying data integrity in critical fields or conducting back-tests to validate the accuracy of a risk model).
- **Exploratory:** Discover hidden patterns or anomalies. Without predefined hypotheses, it explores data freely using visualization and statistical techniques to uncover unknown relationships, structures, clusters, or outliers (for example, applying clustering methods to identify unexpected customer segments or detecting abnormal transactions).

A.3 Tokens Count Analysis

In addition, we have statistics on the distribution of counts across categories, difficulty levels, and insight types. We also computed the average token length of components of the Goal, Insight, and Summary components for the corresponding data. The token length statistics for categories and difficulty levels are illustrated in Figures 7 and 8, while the statistics for insight categories are illustrated in Figure 4 of the main text.

As observed in the category and difficulty-level statistics, the Summary section, which synthesizes all insights, exhibits the longest average token length, whereas Insights show the shortest. When examining different insight categories, the token lengths for Questions remain relatively consistent, while Insights in the Prescriptive category demonstrate the highest token count. This is because the Prescriptive insight type typically requires the provision of extensive optimization measures, recommendations, and strategic suggestions.

B Details of Human Review

We conducted an extensive manual review and annotation in this paper. First, during error analysis of existing datasets, we systematically examined each dataset’s goals, questions, and insights. Next, in constructing the InsightEval dataset, we inspected and annotated preexisting issues and then manually validated the outputs of each generation step. For quality control, a panel of domain experts evaluated the dataset across multiple criteria. Finally, to assess the validity and reliability of our proposed insight metrics, we solicited expert ratings and com-

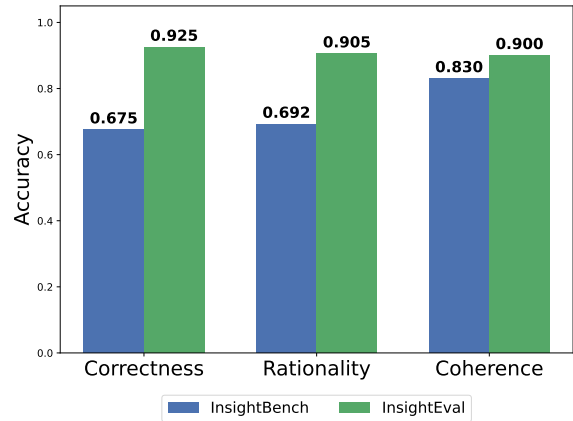


Figure 9: Comparative Benchmark Quality Assessment via LLM Evaluation

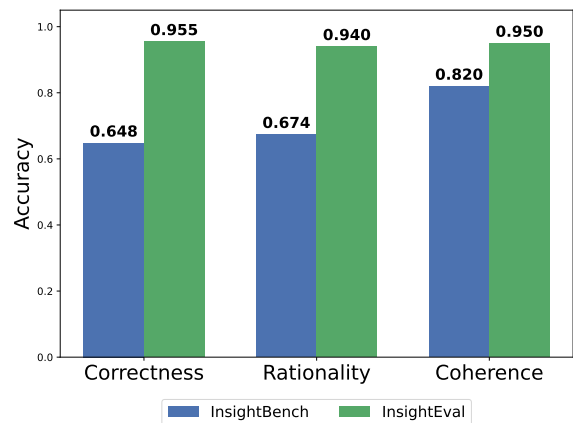


Figure 10: Comparative Benchmark Quality Assessment via Human Evaluation

pared these annotations against our automated measures for rigorous analysis.

C InsightEval vs. InsightBench

To highlight the improvement of our InsightEval dataset, we replicated the LLM-assisted expert annotation and scoring process on InsightBench, following the Data Quality Analysis described in the main text. Compared to the annotation of InsightEval in Table 2 of the main text, the result is presented in the Figure 9 and 10. These comparisons indicate that InsightEval yields enhanced quality and reliability, as corroborated by both human raters and LLM-based assessments.

Since our work is based on an optimized and refined version of the original InsightBench dataset, we have conducted a redundancy analysis on the original InsightBench dataset using the same methodology. The results are presented in Table 6. A comparison with Table 4 in Section 5.2 shows

that our InsightEval achieves varying degrees of improvement across nearly all metrics. This enhancement significantly strengthens the quality and credibility of our dataset.

Type	TC Similarity	Self-BLEU	Distinct-2
Questions	0.0962	0.1746	0.8729
Insights	0.1563	0.2120	0.8216

Table 6: Comparison of average TF-IDF cosine (TC) similarity, Self-BLEU, and Distinct-2 diversity scores between Question and Insight entries in InsightBench.

D Detailed Experimental Setup

D.1 Models Used in this Paper

In the construction of the dataset, the design of the agent framework, and the evaluation process, we employed several widely recognized large language models (LLMs). Detailed information regarding the versions and properties of these LLMs is provided in the Table 7.

Models	Version	Creator
DeepSeek-V3	20250324	deepseek
GPT-4o	20240806	OpenAI
GPT-3.5-Turbo	-	OpenAI
O3-mini	20250416	OpenAI
Claude 3.7 Sonnet	20250219	Anthropic
Gemini 2.5 Pro	20250617	Google

Table 7: Models used in this paper.

D.2 Agent Baselines

We adopted a total of five agent-based baselines, with detailed descriptions provided below.

- **ReAct (Yao et al., 2023)**: ReAct (Reasoning and Acting) is a general paradigm that enhances LLMs’ ability to solve complex tasks by tightly coupling explicit reasoning with external actions. In this framework, the model iteratively produces structured reasoning traces—supporting planning, monitoring, revision, and exception handling—and executes concrete actions that query external resources to gather necessary evidence. In our experiment, we implemented several core tools, including CSV file reading, search functionality, and numerical computation.

- **CodeGen (Majumder et al., 2025)**: The core strategy of CodeGen is to solve a task by generating the complete code in a single pass rather than performing iterative, multi-turn reasoning as in ReAct. Its workflow is as follows: the agent receives a task description, the dataset path, dataset metadata (e.g., column descriptions), and a concrete discovery goal; the prompt context supplies a code demonstration for a similar task (few-shot learning); the model then attempts to produce in one shot all code required to accomplish the task; finally, the generated code is executed and, based on the execution outputs, the model produces a natural-language hypothesis (the final answer) and a concise summary of the workflow. In the experiment, we provided the table and the target as inputs, enabling the code generation model to produce executable code for reading the table and performing analysis, which was then executed. The resulting outputs were subsequently analyzed to derive actionable insights.

- **DeepResearchAgent (Zhang et al., 2025)**: The Deep Researcher Agent is a specialized sub-agent within the AgentOrchestra framework for large-scale, multi-round information retrieval. It can autonomously formulate optimized search queries for complex tasks and perform breadth-first retrieval across multiple search engines. Finally, it synthesizes the extracted insights into a structured research report with explicit source attributions, providing accurate and comprehensive external knowledge to support high-level planning and decision-making. Due to the absence of direct modules or tools for reading CSV tables, the table information was preprocessed and transformed into detailed textual descriptions before being provided to the agent in the experiment.

- **Pandas Agent (LangChain, 2024)**: A data-science agent developed within the LangChain framework that can directly interrogate a Pandas DataFrame. In the experiment, we supply the agent with the table schema and the specified goal. Then it autonomously generates a set of questions, produces corresponding answers and insights, and finally synthesizes these insights into a summary.

- Agent Poirot (Sahu et al., 2025): A popular recent multi-step, multi-round framework for insight generation published in ICLR’25. For each instance, Agent Poirot first extracts the table structure and goal, then generates an initial batch of related questions. To answer these questions, it generates Python code and executes it to obtain table information. The agent answers the corresponding questions and generates insights. Then iteratively formulates extra questions based on earlier outputs, repeating this cycle until a comprehensive set of insights is assembled.

D.3 More Implementation Details.

For experiments involving a single model, to enable the model to process tabular information, we first convert each table into a detailed natural-language description using a Python script in conjunction with LLM-based annotation. The resulting textual representation is then provided to the model as input. Similarly, for both single-agent and multi-agent frameworks, we equip the agents with tools that allow direct reading and querying of CSV-formatted tables.

In addition, for the three primary backbone models used in our experiments, we set the temperature to 0 to ensure deterministic behavior. Regarding output generation, Agent Poirot is configured to run for four iterations, where each iteration produces three new questions and corresponding insights, resulting in a total of 12 generated insights. For fairness and comparability, all baseline methods are required to generate the same number of insights under identical output constraints.

E More Experimental Results

Agents Exhibit a Propensity for Generating Actionable and Exploratory Insights. As illustrated in Figure 11, nearly all agents achieve their highest scores within the Prescriptive, Evaluative, and Exploratory categories. Notably, under both the GPT-4o and Claude-3.7-Sonnet backbones, AgentPoirot attains a peak score of 0.58 in the Prescriptive and Exploratory types. These two types respectively assess an agent’s capacity for recommending executable actions and uncovering latent associations. In contrast to the other categories, the comparatively higher performance in Prescriptive and Exploratory suggests that agents are more adept at formulating practical recommendations and probing potential rules in the dataset.

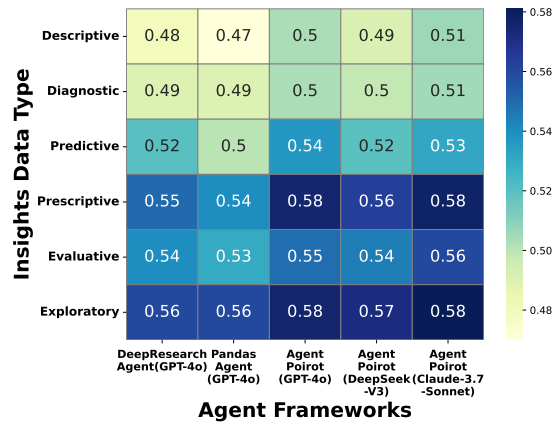


Figure 11: Agent Performance across Insight Types.

Comprehensive Category- and Difficulty-Level Analysis Reveals Superior Performance of Agent Poirot

We add more Experimental Results in Table 8, 9, 10 and 11. Claude-3.7-Sonnet-based Agent Poirot attained the highest overall score among all agents. To further evaluate its performance, we conducted a comprehensive statistical analysis of its scores across various categories and difficulty levels. Across different performance categories, we observed that the agent achieved significantly higher Insight F1 scores in the Asset Management category. In contrast, the Novelty scores were comparatively higher in the Asset & User Management category. Regarding difficulty levels, the agent demonstrated superior performance in Insight F1 scores at Level 1. Similarly, Level 1 yielded higher scores in terms of Novelty. Overall, the scores for the most challenging level, Level 4, were notably lower than those for other levels, suggesting that the difficulty classification is reasonably effective.

F Cases of InsightEval

We illustrate an example of our dataset (data-8) titled Caller Incident Impact Analysis. The metadata is shown in Table 12, which indicate its title, category (Incident Management), role, and difficulty (2). The table of data-8 includes 500 items about simulating ServiceNow incidents. Each item comprises structured fields such as caller id, category, state, opened at, closed at, assigned to, and priority, along with a description of the incident. The Goal of this data is to analyze incident submissions by human callers over time to detect a rising trend relative to peers.

In addition, we illustrate the details of insights

Category	Insights Recall		Insights Precision		Insights F_1	
	ROUGE-1	G-Eval	ROUGE-1	G-Eval	ROUGE-1	G-Eval
Incident Management	0.245	0.513	0.253	0.601	0.249	0.550
Finance Management	0.261	0.572	0.270	0.654	0.265	0.608
Goal Management	0.284	0.601	0.293	0.653	0.288	0.624
User Management	0.283	0.589	0.291	0.645	0.287	0.614
Asset Management	0.288	0.627	0.294	0.695	0.291	0.657
Strategic & Goal Management	0.251	0.446	0.256	0.553	0.253	0.492
Asset & User Management	0.252	0.485	0.258	0.727	0.255	0.582
Finance & User Management	0.287	0.616	0.290	0.673	0.288	0.643
Avg.	0.269	0.556	0.276	0.650	0.272	0.596

Table 8: Performance of categories on Agent Poirot based on Claude-3.7-Sonnet in Insights Level on InsightEval.

Category	Original Novelty Score	New Novelty Score
Incident Management	0.647	0.703
Finance Management	0.732	0.807
Goal Management	0.760	0.802
User Management	0.708	0.740
Asset Management	0.819	0.833
Strategic & Goal Management	0.542	0.646
Asset & User Management	0.875	0.875
Finance & User Management	0.792	0.792
Avg.	0.734	0.775

Table 9: Performance of categories on Agent Poirot based on Claude-3.7-Sonnet in Novelty on InsightEval.

Category	Insights Recall		Insights Precision		Insights F_1	
	ROUGE-1	G-Eval	ROUGE-1	G-Eval	ROUGE-1	G-Eval
Level 1	0.253	0.583	0.255	0.706	0.254	0.633
Level 2	0.273	0.593	0.278	0.655	0.275	0.621
Level 3	0.271	0.583	0.274	0.667	0.272	0.619
Level 4	0.244	0.487	0.250	0.576	0.247	0.525
Avg.	0.260	0.561	0.264	0.651	0.262	0.600

Table 10: Performance of difficulties on Agent Poirot based on Claude-3.7-Sonnet in Insights Level on InsightEval.

and summary in Table 13. Each question is annotated with a data type (Descriptive, Diagnostic, Predictive, Prescriptive, Evaluative, or Exploratory) and a synthesized insight summary.

G Case Study of Comparison between Generated Results and Ground-Truth

To enable a more intuitive comparison between the agent’s output and the ground truth in our dataset, we present an illustrative example in the Table 14 and 16. Notably, the agent generated 12 questions

and corresponding insights, whereas the ground truth only has 10.

According to the evaluation framework, the recall analysis shows that only the I4 ground-truth insight was unsuccessfully captured by the model outputs. For example, I2 ground-truth insight highlighted that the reported number of IT department managers was excessive, which aligns with agent-generated output I2 stating that IT managers faced an overly broad span of control. Both insights reflect a shared observation regarding managerial overload from an organizational average

Category	Original Novelty Score	New Novelty Score
Level 1	0.958	0.958
Level 2	0.756	0.795
Level 3	0.757	0.814
Level 4	0.591	0.659
Avg.	0.766	0.807

Table 11: Performance of difficulties on Agent Poirot based on Claude-3.7-Sonnet in Novelty on InsightEval.

Data Key	Context
Header	Caller Incident Impact Analysis (Data 8)
Category	Incident Management
Role	Resource Manager Analyst
Difficulty	2
Table Description	The dataset comprises 500 entries simulating ServiceNow incidents table, detailing various attributes such as category, state, open and close dates, involved personnel, and incident specifics like description, and priority. It captures incident management activities with fields like 'opened at', 'closed at', 'assigned to', 'short description', and 'priority', reflecting the operational handling and urgency of issues across different locations and categories.
Table Path	./csvs/data-8.csv
User Table Path	null
Table Schema	Column: category (object) missing count: 0 unique count: 5 top5 unique values: ['Database', 'Hardware', 'Inquiry / Help', 'Software', 'Network'] Column: state (object) missing count: 0 unique count: 2 top5 unique values: ['Closed', 'Resolved'] ...
Goal	Examine the distribution of incident tickets submitted by human callers by analyzing the caller id and opened at fields to evaluate frequency over time and identify any caller exhibiting a consistent upward trend in incident submissions relative to peers.

Table 12: Example of metadata in data-8 in InsightEval.

perspective. Besides, in the precision-based analysis, agent-generated outputs I4, I8, and I10 did not correspond to any ground-truth insights, whereas the remaining outputs exhibited strong alignment. For instance, agent-generated output I1 noted a stark 3:1 workload disparity between IT managers Ed Gompf and Mariano Maury, which closely mirrors ground-truth I6.

Overall, our dataset effectively supports agents in uncovering meaningful insights and enables more accurate and comprehensive evaluation of their performance.

H Prompts

Prompt 1, Prompt 2, Prompt 3, Prompt 4, Prompt 5, and Prompt 6 present the detailed prompts for data construction in InsightEval.

Data Key	Context
Question 1 Data Type Insight 1	What is the overall average number of incidents raised by callers over the recent period? Descriptive David Loo has raised a significantly higher number of incidents compared to other callers.
Question 2 Data Type Insight 2	How do the incidents raised by David Loo compare to other agents over the specific same time frame or time period? Descriptive David Loo's incidents are significantly higher and show a linear increasing trend over time compared to other callers.
Question 3 Data Type Insight 3	Are there changes in the categories of incidents raised by David Loo over time? Diagnostic Incidents raised by David Loo are predominantly in the Network category.
Question 4 Data Type Insight 4	What could be potential consequences of not addressing the anomaly or trend of raising tickets by one employee? Predictive There is a continued linear increase in ticket submissions by David Loo.
Question 5 Data Type Insight 5	What are the monthly counts of incident submissions for each caller and how do these counts vary over time? Descriptive Across the months, Bud Richman exhibited the highest submission rate peaking at 17 incidents in June 2023, while David Loo and Don Goodliffe also showed notable submission frequencies, suggesting a diverse pattern of incident reporting among different callers.
Question 6 Data Type Insight 6	What correlations exist between incident categories, priorities, or locations and the observed upward trend in submissions by specific callers? Diagnostic David Loo's submissions highlight a specific trend, as he consistently reports high-priority incidents, indicating a potential area for focused improvement within software services, particularly considering his incidents are concentrated in a critical category.
Question 7 Data Type Insight 7	Based on historical submission patterns, what is the forecasted number of incidents for the caller exhibiting a consistent upward trend over the next quarter? Predictive Notably, the trend suggests a potential increase in workload for the service team, as the caller is projected to submit an average of over 11 incidents per month, indicating rising demand for assistance.
Question 8 Data Type Insight 8	What operational adjustments or resource reallocations can be recommended to address and mitigate the impact of a rising trend in incident submissions by certain callers? Prescriptive Don Goodliffe's incident submissions peaked at 2 incidents on several occasions, suggesting a pattern of recurring issues that, if addressed early, could reduce overall ticket volume significantly.
Question 9 Data Type Insight 9	How complete and reliable are the opened at and caller id data fields for accurately assessing the frequency and trend of incident submissions? Evaluative Each caller, including identifiers such as 'ITIL User' and 'David Loo', submitted incidents uniformly, with a total of 125 incidents recorded per caller, indicating consistent engagement across the board.
Question 10 Data Type Insight 10	Are there any unexpected clusters or patterns in caller incident submissions that suggest emerging trends or anomalies in the data? Exploratory Interestingly, while all callers maintained a submission count of 125, this uniformity suggests potential saturation in incident reporting or coordination among callers, signaling an area for further investigation into their submission behaviors.
Summary	The analysis shows that while several callers (like Bud Richman and Don Goodliffe) exhibit distinct submission patterns, David Loo clearly stands out. His incident submissions are significantly higher, predominantly network-related, and display a consistent, linear upward trend over time relative to his peers. This sharp increase, combined with a focus on high-priority issues, suggests that his submissions could signal emerging service challenges and a potential increase in workload for the service team. Additionally, there appears to be an overall uniformity in submission counts across some callers, indicating possible saturation or coordinated reporting behaviors that may warrant further investigation.

Table 13: Example of insights details and summary in data-8 in InsightEval.

Agent Generated Output	Ground Truth
<p>Q1: How many employees are assigned to each manager within the IT department? I1: The 3:1 workload ratio between IT managers (Ed Gompf: 76 employees vs. Mariano Maury: 25 employees) suggests a potential management bottleneck that could impact employee supervision quality, project oversight, and departmental efficiency.</p>	<p>Q1: How does the number of managers and their distribution across departments affect operational effectiveness? I1: There is a disparity in Managerial Distribution across departments.</p>
<p>Q2: What is the average number of direct reports per manager across all departments compared to the IT department specifically? I2: IT department managers face a management span of control nearly 4 times larger than the organizational average, which likely creates significant challenges in effective supervision, employee development, and workload management that could impact departmental performance and employee satisfaction.</p>	<p>Q2: What is the distribution of reportees in the IT department compared to other departments? I2: There is a disproportionate high number of reportees per manager in the IT Department.</p>
<p>Q3: Who are the specific managers in the IT department, and what is the distribution of direct reports for each of them? I3: The substantial workload imbalance in the IT department, where one manager (Ed Gompf) handles three times more direct reports than the other manager (Mariano Maury), could potentially lead to management inefficiencies, employee oversight issues, and burnout for the overloaded manager.</p>	<p>Q3: Who are the managers with the highest number of reportees? I3: There is a significant disparity among managers in terms of reportee numbers.</p>
<p>Q4: Are there any location-based patterns in the assignment of IT employees to managers that might contribute to the workload imbalance? I4: The IT department shows a 2:1 ratio in geographic management complexity between its two managers, with Ed Gompf overseeing nearly twice as many locations (23) as Mariano Maury (12), while simultaneously managing three times as many employees (76 vs 25).</p>	<p>Q4: What strategic changes can be made to improve managerial distribution and enhance departmental staffing? I4: The current distribution of managers and reportees within the IT department highlights significant imbalances that could impact departmental effectiveness and employee morale. To address the uneven managerial load and enhance operational effectiveness in the IT department, the organization should consider the following strategic changes: 1. Conduct a thorough review of current workload and managerial capacity to identify imbalances and potential areas of overload. 2. Implement a more equitable distribution of reportees among existing managers by reassigning team members based on current workload assessments. 3. Increase the number of managerial positions within the IT department to ensure a more manageable span of control and provide better support and oversight. 4. Offer leadership development programs to existing senior team members to prepare them for managerial roles, expanding the leadership base within the department. 5. Explore the use of technology and automation to reduce administrative burdens on managers, allowing them to focus more on leadership and strategic tasks. 6. Regularly monitor the effects of these changes on departmental performance and employee satisfaction to make further adjustments as necessary.</p>
<p>Q5: What is the ratio of employees to managers in the IT department compared to other departments? I5: The extreme disparity in management workload, with IT managers responsible for over 51 employees each compared to just 9.8 employees per manager in Customer Support, suggests a critical need for reorganization to prevent potential burnout and management inefficiencies in the IT department.</p>	<p>Q5: What is the distribution of reportees among managers in the IT department based on the current data? I5: Ed Gompf has a significantly larger workload compared to Mariano Maury, which may indicate a potential overload on his management responsibilities.</p>

Table 14: Comparison between Agent Poirot (Claude 3.7 Sonnet) generated insights and Ground-Truth in data-27 of InsightEval (1).

Agent Generated Output	Ground Truth
<p>Q6: What is the distribution of employee titles within the IT department, and are certain managers handling more specialized technical roles than others?</p> <p>I6: The IT department has a severe management imbalance with one manager overseeing 74% of all IT staff (76 out of 103 employees), which is more than three times the workload of the other manager (25 employees), potentially creating a bottleneck in supervision and support for technical specialists.</p>	<p>Q6: What factors or patterns in the data explain why some IT managers have significantly more reportees compared to others?</p> <p>I6: Currently, Ed Gompf manages 76 reportees, which is 304% more than Mariano Maury's 25, highlighting a potential risk of burnout and inefficiency in the IT department due to unequal workload distribution.</p>
<p>Q7: What is the exact number of employees assigned to each manager in the IT department, and what percentage of the total IT workforce does each manager oversee?</p> <p>I7: There is a significant workload imbalance in the IT department's management structure, with Ed Gompf responsible for nearly three times as many employees (76) as Mariano Maury (25), which could potentially lead to management inefficiencies and employee oversight challenges.</p>	<p>Q7: Based on the existing schedule and current workload distribution, how might the number of reportees per IT manager evolve in the coming month?</p> <p>I7: With Ed Gompf overseeing three times more users than Mariano Maury, there is a clear risk of management overload, suggesting a need for reassessment of user allocations in the IT department.</p>
<p>Q8: Are there any missing manager assignments within the IT department that could indicate organizational gaps or opportunities for better workload distribution?</p> <p>I8: Although only 1.94% of IT employees lack manager assignments, addressing these gaps is important as unassigned employees may lack proper oversight, potentially creating accountability issues and uneven workload distribution among existing managers in the department.</p>	<p>Q8: What managerial adjustments or workload reallocation strategies can be implemented to mitigate potential overload in the IT department?</p> <p>I8: Currently, Ed Gompf oversees 76 users while Mariano Maury manages only 25, indicating a substantial disparity in workload within the IT department that could lead to management challenges if not addressed.</p>
<p>Q9: Which IT department managers have significantly more direct reports than the department average, and what percentage of the total IT workforce do they manage?</p> <p>I9: The IT department shows a severe workload imbalance with one manager (Ed Gompf) responsible for nearly three-quarters of all IT staff, which could create a management bottleneck and potentially affect employee supervision quality and departmental efficiency.</p>	<p>Q9: How reliable and complete are the 'manager' and 'department' data in identifying imbalances in workload distribution for IT managers?</p> <p>I9: In the IT department, Ed Gompf has significantly more system users assigned to him (76) than Mariano Maury (25), highlighting a notable imbalance in workload distribution that could lead to managerial stress.</p>
<p>Q10: What is the geographical distribution of employees reporting to each IT manager, and are there any managers responsible for teams across multiple locations?</p> <p>I10: 100% of IT managers oversee geographically dispersed teams, with an average of 50.5 employees per manager spread across an average of 17.5 locations, suggesting potential challenges in coordination and team management that could impact IT service delivery.</p>	<p>Q10: Are there any unexpected outliers in the IT department where a manager's reported workload significantly deviates from the norm observed across other departments?</p> <p>I10: This managerial workload disparity in the IT department could lead to inefficiencies and burnout, suggesting a review of user assignments to ensure a more equitable distribution among managers.</p>
<p>Q11: Are there any correlations between the number of direct reports and specific employee titles or roles within the IT department that might indicate specialized workload distribution?</p> <p>I11: There is a significant workload imbalance in the IT department management structure, with Ed Gompf overseeing 76 employees while Mariano Maury manages only 25, suggesting a potential management bottleneck that could affect operational efficiency.</p>	

Table 15: Comparison between Agent Poirot (Claude 3.7 Sonnet) generated insights and Ground-Truth in data-27 of InsightEval (2).

Agent Generated Output	Ground Truth
<p>Q12: What is the distribution of direct reports per manager specifically within the IT department, and which managers have significantly higher or lower numbers compared to the department average?</p> <p>I12: The 3:1 ratio of direct reports between IT managers (76 vs 25) suggests a potential management bottleneck for Ed Gompf, which could impact employee supervision quality and departmental efficiency if not addressed through workload rebalancing.</p>	
<p>Summary: The IT department shows a severe workload imbalance between its two managers, with Ed Gompf overseeing 76 employees (73.79% of IT staff) while Mariano Maury manages only 25 employees (24.27%), creating a 3:1 ratio that likely impacts management effectiveness. IT managers handle significantly higher workloads compared to other departments, with an average of 50.5 direct reports per manager versus the company-wide average of 13.4 reports per manager, representing a 276% higher management burden in IT. The geographical distribution of IT staff creates additional management complexity, with Ed Gompf responsible for employees across 23 different locations while Mariano Maury oversees 12 locations, suggesting that location-based factors contribute to the workload disparity."</p>	<p>Summary: The insights reveal a significant imbalance in the IT department's managerial workload, with some managers overseeing far more system users than others. Specifically, Ed Gompf manages 76 users, while Mariano Maury handles only 25-a disparity that could lead to management overload, inefficiencies, and potential burnout. The recommendations include reviewing and redistributing workloads, increasing managerial positions, developing leadership among senior staff, leveraging technology to reduce administrative burdens, and continuously monitoring the impact of these changes to ensure a more equitable and effective management structure.</p>

Table 16: Comparison between Agent Poirot (Claude 3.7 Sonnet) generated insights and Ground-Truth in data-27 of InsightEval (3).

Prompt 1: Prompt of Goal Refinement.

```
1 Instruction:
2
3 Given the following table description:
4 <description>{table_description}</description>
5
6 Given the following table schema:
7 <context>{table_schema}</context>
8
9 Given the following user table scheme(if has):
10 <context>{user_table_schema}</context>
11
12 Given the following goal:
13 <goal>{goal}</goal>
14
15 Instructions:
16 * Analyze the given GOAL against the provided TABLE DESCRIPTION and TABLE SCHEMA. Perform
17   these checks:
18   1. Relevance: Verify if the goal can be achieved using the table's
19     columns/relationships.
20   2. Feasibility: Determine if the goal is technically achievable (e.g., required
21     operations/joins are possible).
22   3. Clarity: Assess if the goal is specific, unambiguous, and measurable.
23 * You need to refine the given goal by following the above checks concisely and clearly.
24 * Most importantly, your new goal must be strictly enclosed within <goal></goal> tags and
25   give your refined reason within <reason></reason>. Refer to the example response
26   below:
27
28 Example response:
29 <goal>...</goal>
30 <reason>...</reason>
31
32 Response:
```

Prompt 2: Prompt of Question Generation.

```
1 ### Instruction:
2
3 Given the following table description:
4 <description>{table_description}</description>
5
6 Given the following table schema:
7 <context>{table_schema}</context>
8
9 Given the following user table schema(if has):
10 <context>{user_table_schema}</context>
11
12 Given the following goal:
13 <goal>{goal}</goal>
14
15 Given the existing questions and their data type:
16 <questions>{exist_questions}</questions>
17
18 Instructions:
19 * Write a list of new questions to supplement the existing questions. New questions can be
20   solved by the data scientists in your team to explore my table data and reach my goal.
21 * Your questions must be relevant to my goal. You can consider 6 types of questions in
22   data:
23   1. Descriptive: Summarizes what happened by aggregating and visualizing
24     historical data (e.g., plotting monthly portfolio returns).
25   2. Diagnostic: Explains why it happened by identifying correlations, patterns,
26     and root causes (e.g., segmenting losses by asset class to find drivers).
27   3. Predictive: Forecasts what is likely to happen using statistical models on
28     past trends (e.g., forecasting next quarter's default risk based on credit scores).
29   4. Prescriptive: Recommends actions to take to optimize outcomes or mitigate
30     risks (e.g., suggesting portfolio rebalancing to reduce predicted volatility).
31   5. Evaluative: Assesses the quality, reliability, and robustness of the data
32     and existing analyses/models (e.g., checking data completeness for key fields or
33     backtesting a risk model's accuracy).
34   6. Exploratory: Discovers hidden patterns, relationships, or anomalies without
35     predefined hypotheses (e.g., using clustering to find unexpected customer segments or
36     identifying outlier trades).
37 * Considering the existing questions and their data types. Make sure that there is at
38   least one question for each type.
39 * You must ask the right questions to surface anything interesting (trends, anomalies,
40   etc.)
41 * Make sure these questions can realistically be answered based on the table description,
42   table schema, and user table schema(if has).
43 * The insights that your team will extract will be used to generate a report.
44 * Each question should only have one part, that is, a single '?' at the end, which only
45   requires a single answer.
46 * Do not number the questions.
47 * You should produce {max_questions} questions to supply. Stop generating after that.
48 * Most importantly, each question type and question must be enclosed within
49   <data_type></data_type> and <question></question> tags. Refer to the example response
50   below:
51
52 Example response:
53 <data_type>Descriptive</data_type>
54 <question>What is the distribution of the customers based on their age?</question>
55
56 ### Response:
```

Prompt 3: Prompt of Code Generation.

```
1 ### Instruction:
2
3 Given the goal:\n
4 {goal}
5
6 Given the schema:\n
7 {schema}
8
9 Given the data path:\n
10 {database_path}
11
12 Give me the Python code required to answer this question "{question}" and put a comment on
13   top of each variable.\n\n
14 Instructions:
15 * Make a single code block for starting with ```python
16 * Do not produce code blocks for languages other than Python.
17 * Import pandas as pd, and numpy as np at the beginning.
18 * Save a stats JSON file that stores the data.
19 * Each JSON file must have a "name", "description", and "value" field that describes the
20   data.
21 * If the content of the JSON file is getting too long, truncate the unnecessary parts.
22 * End your code with ```.
23 Output code:\n
```

Prompt 4: Prompt of Question Answering.

```
1 ### Instruction:
2 You are trying to answer a question based on the information provided.
3
4 Given the following dataset schema:
5 <schema>{schema}</schema>
6
7 Given the goal:
8 <goal>{goal}</goal>
9
10 Given the question:
11 <question>{question}</question>
12
13 Given the data information:
14 <data>{data}</data>
15
16 Instructions:
17 * Based on the data and other information provided above, write an answer to the question
18   enclosed with <question></question> tags.
19 * The answer should be a single sentence, but it should not be too high-level and should
20   include the key details from the justification.
21 * Write your answer in HTML-like tags, enclosing the answer between <answer></answer>
22   tags, followed by a justification between <justification></justification> tags.
23 * Refer to the following example response for the format of the answer and justification.
24
25 Example response:
26 <answer>This is a sample answer</answer>
27 <justification>This is a sample justification</justification>
28
29 ### Response:
```

Prompt 5: Prompt of Insight Generation.

```
1 ### Instruction:
2 You are trying to produce an insight based on information provided .
3
4 Given the following dataset schema:
5 <schema>{schema}</schema>
6
7 Given the goal:
8 <goal>{goal}</goal>
9
10 Given the question:
11 <question>{question}</question>
12
13 Given the answer:
14 < answer>{answer}</answer >
15
16 Instructions:
17 * Produce an insight that we should look into to explore my data and reach my goal.
18 * Write your insight in HTML-like tags, enclosing the insight between <insight></insight>
19   tags.
20 * The insight should be something interesting and grounded based on the question, answer,
21   goal, and the dataset schema, something that would be interesting.
22 * The insight should be as quantitative as possible and informative and non-trivial and
23   concise.
24 * The insight should be a meaningful conclusion that can be acquired from the analysis in
25   laymans terms
26
27 Example response:
28 <insight>This is a sample insight</insight>
29
30 ### Response:
```

Prompt 6: Prompt of Summary Synthesis.

```
1 ### Instruction:
2 You are trying to answer a question based on information provided by a data scientist.
3
4 Given the goal:
5 <goal>{goal}</goal>
6
7 Given the insights list:
8 <insights>{insights}</insights>
9
10 Instructions:
11 * Given a goal, and all the insights from the above list. Give your concise summary of the
12   above insights.
13 * The summary should be more about the highlights of the findings.
14 * Your output summary should within this tag <summary></summary>.
15
16 ### Response:
```