

JarviX: A LLM No code Platform for Tabular Data Analysis and Optimization

Shang-Ching Liu^{1,3,†}, ShengKun Wang^{2,3,†}, Wenqi Lin³, Chung-Wei Hsiung³,
Yi-Chen Hsieh³, Yu-Ping Cheng³, Sian-Hong Luo³, Tsungyao Chang³, Jianwei Zhang^{*1}

¹Department of Computer Science, University of Hamburg

²Department of Computer Science, Virginia Tech

³Synergies Intelligent Systems, Inc.

`jianwei.zhang@uni-hamburg.de`

Abstract

In this study, we introduce JarviX, a sophisticated data analytics framework. JarviX is designed to employ Large Language Models (LLMs) to facilitate an automated guide and execute high-precision data analyzes on tabular datasets. This framework emphasizes the significance of varying column types, capitalizing on state-of-the-art LLMs to generate concise data insight summaries, propose relevant analysis inquiries, visualize data effectively, and provide comprehensive explanations for results drawn from an extensive data analysis pipeline. Moreover, JarviX incorporates an automated machine learning (AutoML) pipeline for predictive modeling. This integration forms a comprehensive and automated optimization cycle, which proves particularly advantageous for optimizing machine configuration. The efficacy and adaptability of JarviX are substantiated through a series of practical use case studies.

1 Introduction

Although the predominant focus of contemporary research on large language models is the evaluation of various tasks (Liang et al., 2022; Zhao et al., 2023), there is a noticeable lack of academic resources that provide structured guidelines and frameworks for downstream applications. Tabular data analysis, as an important application task of LLMs, has always faced challenges related to the precision of mathematical calculations. Despite its ability to address complex high school math problems and participate in advanced mathematical discussions, advanced models such as GPT-4 are not yet on par with expert level performance (Bubeck et al., 2023). They are prone to basic errors and can sometimes produce incoherent outputs.

[†]These authors contributed equally to this work.

This may stem from the fact that autoregressive models lack self-correction mechanisms when generating solutions (Shen et al., 2021). This paper introduces a thorough approach towards employing LLMs for tabular data analysis tasks, specifically aiming to equip nonspecialists with the ability to engage in advanced data analytics using LLMs within a rule-based system. Although LLMs have proven to be potent in data processing (Zhao et al., 2023), their application in guiding users through rule-based systems to intuitively create data visualizations, synthesize statistical insights, and provide context-aware explanations is significantly under-explored.

LLM guides users through a rule-based system in JarviX, which enables data visualization and statistical analysis. It leverages a vectorized domain knowledge repository to provide relevant explanations for each visualization and suggests further exploration directions (Feng et al., 2023). Users can generate additional exploratory charts and insights through text or voice input, facilitated by Whisper.¹ These insights are processed by the Vicuña model² through prompts fine-tuned by GPT-4³. This approach culminates in a comprehensive report encapsulating all the insights and analytical processes, serving as a thorough guide for users and a blueprint for future analysis.

Furthermore, this study explores the integration of H2O-AutoML-customized AutoML pipelines (H2O.ai, 2023) into this process. The use of AutoML is demonstrated to identify the best targets with respect to other data columns and to build specific models to optimize results for various objectives, such as optimal factory configurations (He et al., 2021).

¹<https://github.com/openai/whisper>

²<https://lmsys.org/blog/2023-03-30-vicuna/>

³<https://openai.com/research/gpt-4>

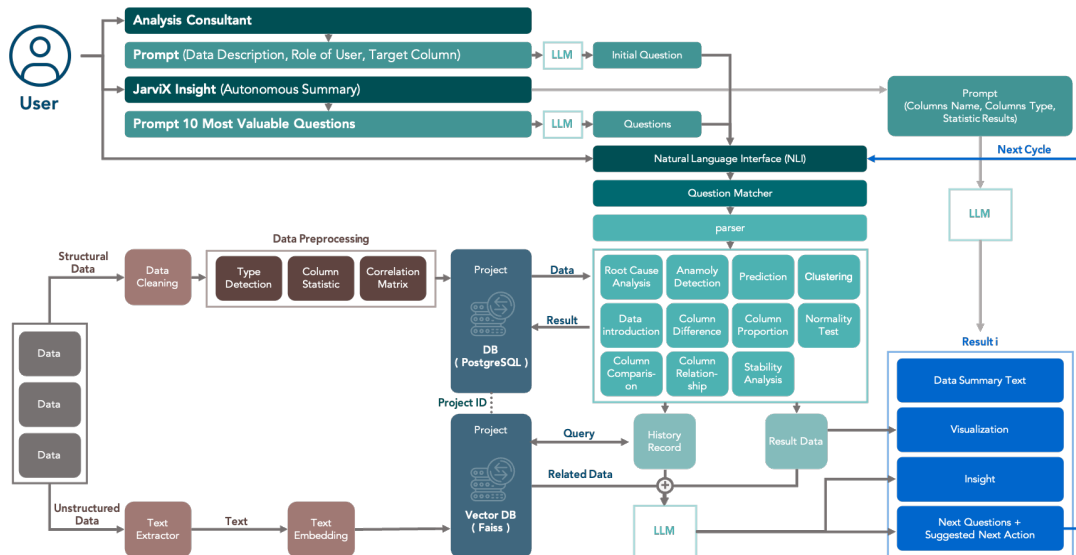


Figure 1: JarviX system overview

The primary objective of this study is to empower users with the knowledge and tools necessary to harness the power of LLM for rule-based data analytics by fine-tuning (Chung et al., 2022) and AutoML. The paper concludes by underlining the potential of this approach in democratizing data analytics, thereby fostering more strategic and informed decision-making.

2 Related Work

2.1 Natural Language Interfaces for Data analysis

Natural Language Interfaces have recently garnered attention and integration into various commercial data analysis and visualization software, such as IBM Watson Analytics (IBM, 2023), Microsoft Power BI (Microsoft, 2023), Tableau (Salesforce, 2023), and Google Spreadsheet (Google, 2023). Despite initial limitations, such as confining natural language interactions to data queries and standard chart types, the approach is evolving. Current methods of Natural Language Processing (NLP) incorporate heuristic algorithms, rule-based systems, and probabilistic grammar-based approaches, each with their respective challenges and trade-offs in accuracy, flexibility, and computational resources (Miwa and Bansal, 2016; Voigt et al., 2021; Satyanarayan et al., 2016).

2.2 Utilization of LLMs in Advanced Data Analysis

Rajkumar’s performance evaluation of LLM on Text2SQL (Rajkumar et al., 2022) was a significant

development. Despite the considerable progress, including Sun et al.’s (Sun et al., 2023) Text2SQL method achieving 77.3% accuracy on the Spider benchmark, constructing a seamless pipeline is still challenging. The optimization strategy of Hu et al. (Hu et al., 2023) and the question refinement strategy of Guo et al. (Guo et al., 2023) show further improvements. The Maddigan et al. system (Maddigan and Susnjak, 2023) underscores the importance of visualization post-SQL generation. However, the demand for practical solutions for real-world applications remains, prompting the development of the JarviX platform. It bridges the gap, offering higher-level APIs and integrating LLMs for a comprehensive solution for advanced data analysis.

2.3 External Knowledge Integration

Currently, LLMs are confronted with two issues: privacy implications and the obsolescence of training data. Utilizing user-interaction data for further training in online LLM applications can potentially jeopardize security. Additionally, the significant costs associated with retraining can make LLMs outdated over time. LangChain (Chase, 2022) provides an innovative solution that continuously embeds the latest data and retrieves relevant information from its database, consequently generating current responses while preserving privacy. Moreover, the introduction of the llama_index (Liu, 2022) proposes a more structured approach for embedding levels to retrieve query-related information, such as identifying the latest facts or providing relational

data. This method improves the precision of LLM responses.

3 Overview

JarviX is a no-code platform for efficient analysis and optimization of tabular data, handling both structured and unstructured types, as illustrated in Figure 1. For structured data (e.g., csv files, data frames), it performs preliminary processing tasks including data type detection, statistical computation, and correlation analysis, storing the results in a Postgres database. Unstructured data (e.g., text files, audio files) is managed through text extraction and embedding, followed by storage in a vectorized database, such as Elastic Search.

Users can interact with the platform through three key features: JarviX Insight, Natural Language Interfaces, and JarviX Guidance. JarviX Insight collects structured data information such as column names, types, and statistical data, and employs a LLM to generate a data summary report, providing users with an understanding of their data and identifying key questions.

The Natural Language Interfaces feature accommodates user queries about their datasets, either voice-to-text or typed, and translates these queries into a rule-based system via a fine-tuned LLM. This delivers relevant data visualizations, explanations, and follow-up question suggestions.

Lastly, JarviX Guidance assists users through a step-by-step data analysis process. It takes into account the user's understanding of the datasets, their role, the specific dataset, and target column they wish to analyze. Using this information, JarviX anticipates the questions a user might want to address first and commences the result generation process. It also provides an appropriate endpoint for analysis for each user. All stages of the analysis are recorded, including the middle *result_i*, and compiled into a comprehensive report that users can save and share.

4 System Break Down

4.1 Data Input Methods

JarviX offers users three methods for uploading structured data: via SFTP, database connections, or direct CSV file uploads. For unstructured data, the platform currently supports only file uploads.

JarviX integrates a data cleaning interface with automated functions, enabling users to efficiently perform data cleaning. Then, the system initiates

data pre-processing, which includes automatic type detection, column statistics computation, and calculation of the correlation matrix between columns. It is worth noting that these tasks—column statistical computation and correlation matrix calculation—are executed asynchronously, ensuring that user progress isn't hindered.

In handling unstructured data, we leverage various connectors in the llama hub (Zhang, 2023) to perform text extraction. The extracted data is stored in the vector database using Faiss (Meta, 2023) and assigned the same project ID as the structured database. This cohesive data management strategy ensures seamless integration and retrieval of both structured and unstructured data.

4.2 JarviX Insight

The JarviX Insight feature facilitates autonomous report generation, enabling users to comprehend data more effectively and gain insight into potential subsequent questions, as illustrated in Figure 1 and shown in Figure 3. Upon activation of the JarviX Insight function, two distinct processes occur. Initially, a prompt with preprocessed data is used to determine the nature of the data, which subsequently assists in the creation of a data summary text. Currently, the LLM is used to generate the ten most pertinent questions. These questions serve as a foundation for generating a variety of potential visualization results. By integrating these elements, a comprehensive data summary report is crafted.

4.3 Question Matcher

The Question Matcher is the key module that links questions from a natural language interface to their corresponding modules using SQL matching. This process relies on identifying three types of keywords: 1) column name-related terms, 2) restriction-related phrases (e.g. "top ten"), and 3) algorithm or module keywords. Once these keywords are identified, the module begins to merge the specific restrictions associated with each column into a unified combination. This combination is then matched with the algorithm or module indicated by the third type of keyword. More details on Question Matcher are contained in the Appendix A.

4.4 Analysis Consultant

The Analysis Consultant is designed for users who have an initial comprehension of their data and express interest in exploring specific columns, as

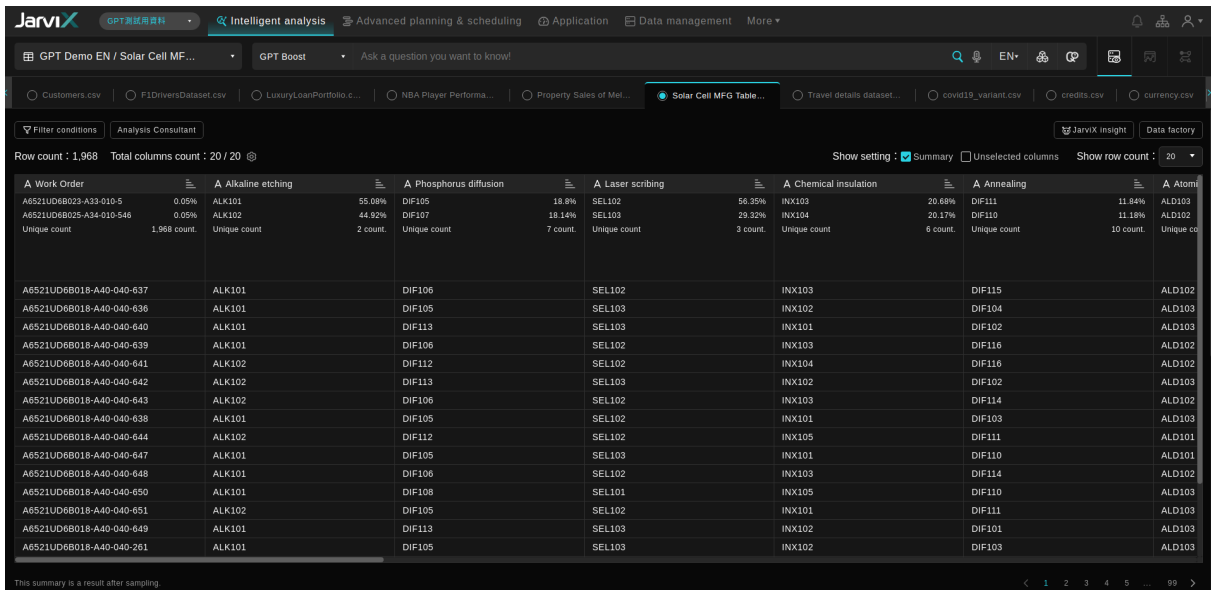


Figure 2: Main page

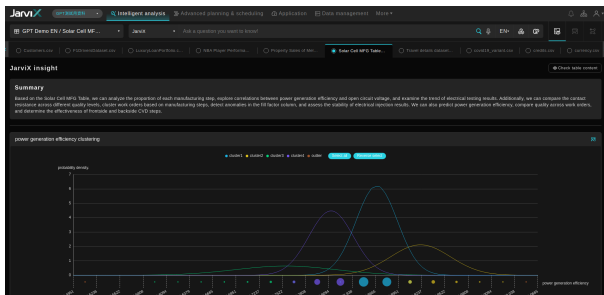


Figure 3: JarviX insight

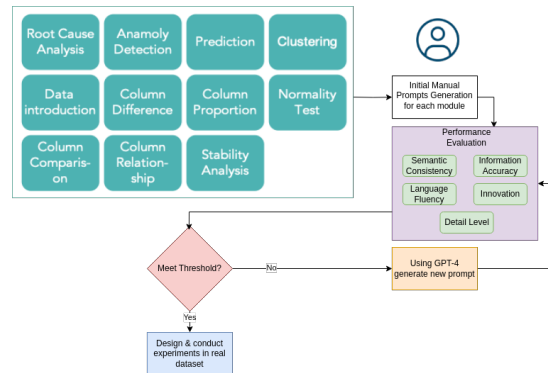


Figure 4: Prompt optimization process

depicted in Figure 1. The process begins with the setup of the analysis parameters based on previously outlined criteria. Subsequently, the LLM formulates the first query. The Consultant then generates comprehensive results that include visualizations, insights with supportive explanations, and prompts for potential follow-up queries from the users. A crucial aspect of this process is the incorporation of professional knowledge into the insights, providing not only a fundamental explanation of the visualizations but also integrating general knowledge and background understanding into the explanation. If the analysis process is deemed comprehensive, the Consultant may propose generating a report.

LLMs put forth subsequent analytical queries, facilitating users in delving deeper into their target column's data features. The formulation of questions is based on the preceding results of the analysis and the roles selected by the user. This feature equips people who lack in-depth data analysis

knowledge to yield thorough and dependable interpretations of their data. A step-by-step description of this process will be exhibited in Session 5.

4.5 LLM, Prompt Engineering

In our experiments, we leverage the vicuna-13b-1.1-gptq-4bit-128g (TheBloke, 2023) as our base LLM, and optimize the prompts with the advanced language understanding capabilities of GPT-4 to achieve better results.

Our approach, based on prompt engineering, can be viewed as a two-stage process, as detailed in Figure 4. First, we manually generate prompts, which require a deep understanding of the task module and a clear view of the structure and goal of the task. The resultant prompts are both specific and concise.

Post-initial generation, we instigate a feedback loop to optimize the prompts. Every prompt is




Figure 5: Question matcher

fed into the Vicuña model to generate respective outputs, which are compared to the expected results to derive a performance metric. If a prompt falls short of our performance benchmark, it's replaced with a new one generated by GPT-4.

We persist with this prompt optimization until all prompts meet our performance criteria. To corroborate our approach, we undertake a range of experiments spanning tasks like normality tests, forecasting, comparisons, root cause analysis, anomaly detection, and relationship extraction.

The primary aim of our experiments is to juxtapose the performance of our approach against existing prompt engineering techniques. The results, presented through multiple examples from our experiments, attest to the strength of our methodology. Taking advantage of the interaction between prompts and feedback loops, the Vicuña model yields rich and insightful responses to complex data-related queries.

5 Case Study

In this section, we demonstrate two separate use cases: 1) Utilizing JarviX Insight for custom analysis with data matching, and 2) Exploring JarviX's guidance use cases.

5.1 Case 1: JarviX Insight and Solar Cell Manufacturing

We present a case that explores solar cell manufacturing data and demonstrates how to increase efficiency using JarviX Insight.

As shown in Figure 2, users begin by using JarviX Insight to generate a report for general understanding of the data set. If users are unfamiliar with the data, the JarviX Insight function can provide a general report that answers two questions: 1) What is the subject of these data? 2) What are the most valuable queries that can be made using this dataset?

Upon gaining insights that quality could be a potential area for enhancement, users can utilize the "Question Matching" feature. This function facilitates the formation of general queries, such as "What is the difference between high quality and low quality" As illustrated in Figure 5, the "Question Matcher" successfully translates user inputs into keywords recognizable by a rule-based system.

The visualization results of different algorithms indicate the key differences between high and low quality. The visualization includes straightforward insights, suggested questions, actions, and the main diagram.

Our AutoML pipeline simplifies the process of training a machine learning model. Users simply define the data source, dataset, and target column,




Figure 6: Simulation

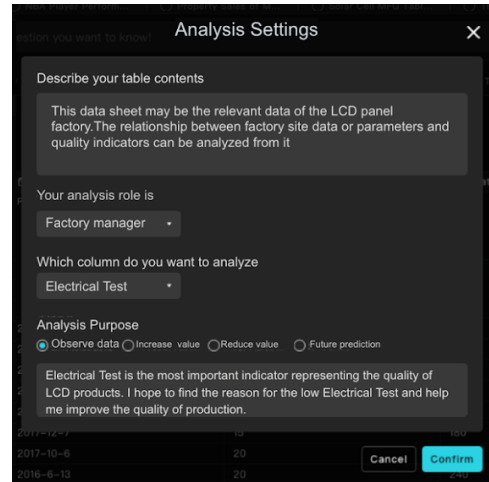


Figure 7: Settings for the analysis

as well as the performance metric for optimization (such as MAE, MSE, or RMSE). Though the training strategy’s precision may impact the model’s training time, once all parameters are set, the model can be generated. Once the model is established, users can explore optimal settings through the simulation panel, as depicted in Figure 6. This panel allows users to identify optimal configurations within the defined range. Importantly, the model is designed to progressively refine itself with the influx of new streaming data. This dynamic adaptation promises improved outcomes over time as settings are intelligently adjusted in response to the evolving data.

This process provides an ideal optimization cycle for customers. In this case, the solar panel manufacturer increased efficiency by 10% using this optimization cycle.

5.2 Case 2: JarviX Guidance, An Analysis of LCD Factory Data

If users are new to JarviX, understanding its functionality or learning how to analyze data might be challenging. To address this, our analysis consultant is available to guide you through the process. In this case study, our focus is on interpreting a dataframe relevant to an LCD panel factory.

Setting up the system properly is paramount to ensure it identifies the data pertinent to the user. We commence by describing the content of the data table, followed by outlining our analysis objectives and roles, as depicted in Figure 7.

As an initial step, JarviX recommends the appropriate analysis that users should consider. It’s often challenging for users to extract vital information

from the analysis process when faced with varied data. Therefore, we inform our model about the analysis approach, so it can recommend suitable subsequent analyses based on users’ requirements.

Upon starting the analysis, JarviX assists users in interpreting the results and guides them to the next steps.

Through our differential analysis, we determined that the electrical test performance heavily depends on the stability of ambient humidity. A list on the left displays the significant factors that influence the differences between high and low electrical tests. In particular, humidity is the top factor, indicating that humidity differences significantly affect the performance of electrical tests.

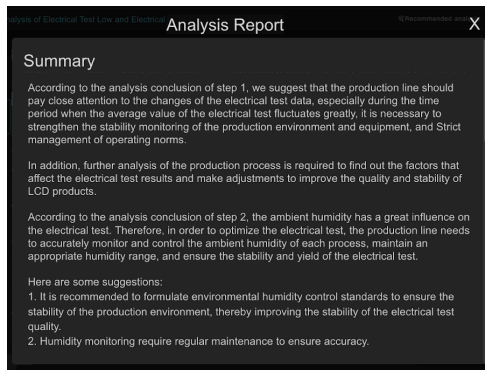
At this stage, the system suggests the production of a summary report based on the previous analysis. JarviX will first show a summary suggestion, then recapitulate the previous analysis steps as Figure 8. The analysis consultant’s guidance enables the user to obtain valuable analysis results, which can help optimize company strategy or uncover potential business value.

6 Conclusion

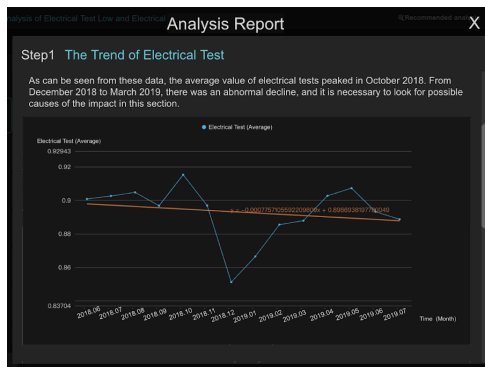
JarviX, by integrating LLM and AutoML technologies, presents a unique and all-encompassing approach for the analysis of tabular data. The system integrates non-structural data to generate profound insights, employing LLM to aid users in their data exploration endeavors.

7 Future Work

The flexibility and adaptability of the JarviX platform offer several avenues for future improvements.



(a) Summary text



(b) Summary figure

Figure 8: Analysis summary

In particular, there are opportunities to fine-tune the LLM to improve personalized recommendations, extend the range of accepted data types and query categories, and improve user interface design for a better user experience.

Ethical Considerations and Limitations

JarviX formulates responses influenced by the user-provided context. Biased results may arise if the context involves biases related to aspects such as the location or language of the user (Hadi et al., 2023). For example, given that JarviX currently only supports processing and analysis in English and Chinese, it might yield biased answers when inquiries about a specific culture or religion are presented, especially if JarviX lacks adequate training in that particular cultural or religious context, due to its confined knowledge. Also, JarviX is designed to recognize only plain text information and cannot identify multimodal tabular data, such as financial statements or instructional videos.

Acknowledgements

This research is supported by Synergies Intelligent Systems, Inc. and partially supported by

the DFG/NSFC Collaborative Project “Crossmodal Learning - Adaptivity, Prediction and Interaction” SFB/TRR169.

References

- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Harrison Chase. 2022. [LangChain](#).
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Yingchaojie Feng, Xingbo Wang, Bo Pan, Kam Kwai Wong, Yi Ren, Shi Liu, Zihan Yan, Yuxin Ma, Huamin Qu, and Wei Chen. 2023. Xnli: Explaining and diagnosing nli-based visual data analysis. *IEEE Transactions on Visualization and Computer Graphics*.
- Google. 2023. [Google sheets](#).
- Chunxi Guo, Zhiliang Tian, Jintao Tang, Shasha Li, Zhihua Wen, Kaixuan Wang, and Ting Wang. 2023. [Retrieval-augmented gpt-3.5-based text-to-sql framework with sample-aware prompting and dynamic revision chain](#).
- H2O.ai. 2023. [h2o-automl](#). 3.42.0.
- Muhammad Usman Hadi, R Qureshi, A Shah, M Irfan, A Zafar, MB Shaikh, N Akhtar, J Wu, and S Mirjalili. 2023. A survey on large language models: Applications, challenges, limitations, and practical usage. *TechRxiv*.
- Xin He, Kaiyong Zhao, and Xiaowen Chu. 2021. Automl: A survey of the state-of-the-art. *Knowledge-Based Systems*, 212:106622.
- Chenxu Hu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Zhao, and Hang Zhao. 2023. [Chatdb: Augmenting llms with databases as their symbolic memory](#).
- IBM. 2023. [Ibm watson analytics](#). 3.5.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekogul, Mirac Suzgun, Nathan Kim,

- Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2022. [Holistic evaluation of language models](#).
- Jerry Liu. 2022. [LlamaIndex](#).
- Paula Maddigan and Teo Susnjak. 2023. [Chat2vis: Generating data visualisations via natural language using chatgpt, codex and gpt-3 large language models](#).
- Meta. 2023. GitHub - facebookresearch/faiss: A library for efficient similarity search and clustering of dense vectors. — [github.com](https://github.com/facebookresearch/faiss). <https://github.com/facebookresearch/faiss>. [Accessed 17-Jul-2023].
- Microsoft. 2023. [Microsoft power bi](#).
- Makoto Miwa and Mohit Bansal. 2016. End-to-end relation extraction using lstms on sequences and tree structures. *arXiv preprint arXiv:1601.00770*.
- Nitarshan Rajkumar, Raymond Li, and Dzmitry Bahdanau. 2022. [Evaluating the text-to-sql capabilities of large language models](#).
- Salesforce. 2023. [Tableau](#).
- Arvind Satyanarayan, Dominik Moritz, Kanit Wongsuphasawat, and Jeffrey Heer. 2016. Vega-lite: A grammar of interactive graphics. *IEEE transactions on visualization and computer graphics*, 23(1):341–350.
- Jianhao Shen, Yichun Yin, Lin Li, Lifeng Shang, Xin Jiang, Ming Zhang, and Qun Liu. 2021. [Generate & rank: A multi-task framework for math word problems](#).
- Ruoxi Sun, Sercan O. Arik, Hootan Nakhost, Hanjun Dai, Rajarishi Sinha, Pengcheng Yin, and Tomas Pfister. 2023. [Sql-palm: Improved large language model adaptation for text-to-sql](#).
- TheBloke. 2023. vicuna-13b-1.1-gptq-4bit-128g model. <https://huggingface.co/TheBloke/vicuna-13B-1.1-GPTQ-4bit-128g?doi=true>. Accessed: 2023-07-20.
- Henrik Voigt, Monique Meuschke, Kai Lawonn, and Sina Zarrieß. 2021. [Challenges in designing natural language interfaces for complex visual models](#). In *Proceedings of the First Workshop on Bridging Human-Computer Interaction and Natural Language Processing*, pages 66–73, Online. Association for Computational Linguistics.
- Jesse Zhang. 2023. [Emptycrown/llama-hub: A library of data loaders for llms made by the community – to be used with gpt index and/or langchain](#).
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#).

TABLE I : Evaluation result for Question Matching. Top1 indicates the instances where the correct result was identified as the top most result. Top3 represents the cases where the correct result was found within the top three results.

Data source	Column Name		Intention		Restriction	
	Top1	Top3	Top1	Top3	Top1	Top3
Manufacture	72.0	83.3	74.0	82.7	64.7	72.0
Sport	73.3	88.3	75.0	90.8	65.8	76.7
Sales	70.7	82.7	77.3	86.7	67.3	78.7
Food	69.2	88.3	75.8	90.8	70.0	77.5
Health Care	65.0	74.2	73.3	85.0	67.5	74.2
Banking	81.7	93.3	79.2	91.2	66.7	74.2

A Experiment Results

A.1 Dataset

We conducted evaluations on JarviX using a variety of tabular datasets⁴ sourced from open source collections, covering different fields. A set of 10 manually crafted questions was complemented by an additional 20 generated by GPT-4 for each data set. To ensure relevance and meaningfulness, the prompts were designed with the phrase “assuming that you are a professional data analyst in this field”, tailoring the questions generated to the specific industry. The test data were then classified into six distinct industry-based segments.

A.2 Results

The TABLE I presents the performance results of the question matching evaluation, focusing on three specific aspects: column name, intention, and restriction detection. Each aspect of the evaluation is thoroughly examined, with results meticulously tabulated to offer a comprehensive understanding of the system’s performance across the different dimensions. JarviX demonstrates proficient recognition of individual columns. However, when faced with questions that encompass multiple columns, there is a possibility that it might not fully recognize all of them. The eleven intentions that JarviX is capable of executing are illustrated in Figure 4. In addition, JarviX is equipped to identify the following specific restrictions: *Average, Median, Sum, Greater than, Equal to, Less than, Plus, Minus, Multiply, Divide, Top, Last, Maximum, Minimum*. On the basis of our experimental findings, JarviX shows enhanced performance as user queries exhibit clearer intent. However, when faced with ambiguous queries, JarviX is prone to over-identifying or under-identifying terms.

⁴<https://reurl.cc/jv693q>