

RETQA: A Large-Scale Open-Domain Tabular Question Answering Dataset for Real Estate Sector

Zhensheng Wang¹, Wenmian Yang^{2*}, Kun Zhou¹, Yiquan Zhang³, Weijia Jia^{2,4*}

¹School of Artificial Intelligence, Beijing Normal University, Beijing 100875, PR China

²Institute of Artificial Intelligence and Future Networks, Beijing Normal University, Zhuhai 519087, PR China

³Elmleaf Ltd., Shanghai 200082, PR China

⁴BNU-UIC Institute of Artificial Intelligence and Future Networks, Beijing Normal University (Zhuhai), Guangdong Key Lab of AI and Multi-Modal Data Processing, BNU-HKBU United International College, Zhuhai, Guang Dong, PR China
{jensenwang,zhoukun}@mail.bnu.edu.cn, {wenmianyang,jiawj}@bnu.edu.cn, zhangyq987@hotmail.com

Abstract

The real estate market relies heavily on structured data, such as property details, market trends, and price fluctuations. However, the lack of specialized Tabular Question Answering datasets in this domain limits the development of automated question-answering systems. To fill this gap, we introduce RETQA, the first large-scale open-domain Chinese Tabular Question Answering dataset for Real Estate. RETQA comprises 4,932 tables and 20,762 question-answer pairs across 16 sub-fields within three major domains: property information, real estate company finance information and land auction information. Compared with existing tabular question answering datasets, RETQA poses greater challenges due to three key factors: long-table structures, open-domain retrieval, and multi-domain queries. To tackle these challenges, we propose the SLUTQA framework, which integrates large language models with spoken language understanding tasks to enhance retrieval and answering accuracy. Extensive experiments demonstrate that SLUTQA significantly improves the performance of large language models on RETQA by in-context learning. RETQA and SLUTQA provide essential resources for advancing tabular question answering research in the real estate domain, addressing critical challenges in open-domain and long-table question-answering.

Code — <https://github.com/jensenw1/RETQA>

Introduction

With rapid advancements in artificial intelligence and natural language processing, Tabular Question Answering (TQA) has garnered attention for its ability to extract accurate answers from structured data across domains like finance and healthcare (Zhu et al. 2024). Despite progress, specialized datasets for the real estate domain are scarce, limiting research and practical applications. The real estate market relies on structured data, such as property details and market trends, which are crucial for stakeholders like homebuyers and investors. For instance, homebuyers can make informed decisions by querying historical transaction prices and property details, while investors can assess the viability

of investments by querying land information, asset status, and the developer’s transaction history. However, the lack of tailored TQA datasets in this field hampers the development of automated QA systems.

To address this gap, we introduce RETQA, the first large-scale open-domain Chinese TQA dataset specifically designed for the Real Estate sector. RETQA is built from publicly available real estate data and comprises 4,932 tables and 20,762 QA pairs spanning 16 sub-fields within three major domains: property information, real estate company finance, and land auction information. Given the complex and data-intensive nature of the real estate market, RETQA is crafted to generate queries targeting longer tables, with an average of 252.9 rows per query-related table. Moreover, queries in RETQA may related to more than one table. The above design captures the real estate domain’s complexity, making RETQA a challenging TQA dataset.

The open-domain nature of RETQA requires models to retrieve relevant tables from the entire dataset, rather than being directly provided, effectively mirroring real-world scenarios. However, this process makes open-domain TQA more challenging than closed-domain, as retrieval accuracy directly affects overall TQA performance. To facilitate retrieval, RETQA assigns a summary caption to each table. Furthermore, RETQA integrates the labels of Spoken Language Understanding (SLU) for each query. These SLU labels, including intent and slot labels, commonly employed in task-oriented dialogue systems (Qin et al. 2021; Cheng, Yang, and Jia 2023; Zhu et al. 2024; Qin et al. 2024), are instrumental in discerning user query intent and extracting pertinent details. As depicted in Figure 1, SLU labels enable more accurate parsing of user intent and relevant information, thereby enabling more precise retrieval and accurate answering. For each query, RETQA provides intent labels, slot labels, table captions, and answers in Markdown, SQL-style, and natural language formats. To the best of our knowledge, RETQA is the first TQA dataset to integrate SLU labels.

RETQA is created through four key steps: First, we compile a dataset of 4,932 real estate-related tables from eight major Chinese cities—Beijing, Shanghai, Guangdong, Shenzhen, Suzhou, Hangzhou, Nanjing, and

*Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

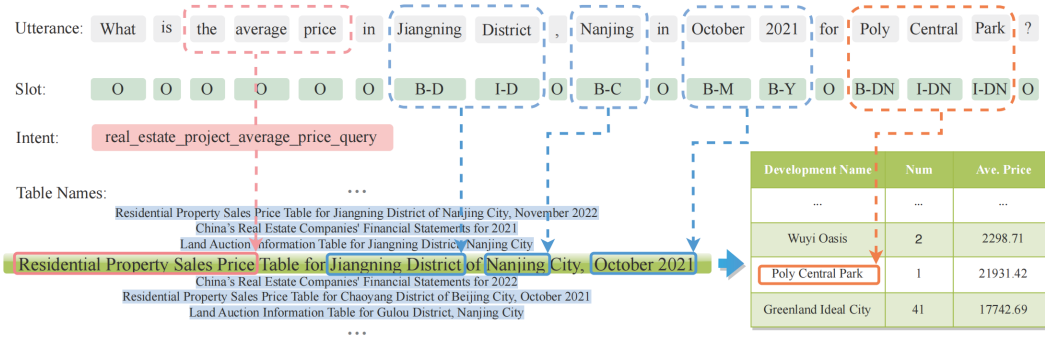


Figure 1: Example of SLU labels and the relation to open-domain TQA, where “D” represents district, “C” represents city, “M” represents month, “Y” represents year, and “DN” represents development name.

Wuhan—spanning the years 2019 to 2022. These tables include property information, real estate company finances, and land auction data. The publicly available data is meticulously cleaned and organized, and each table is assigned a summary caption to aid retrieval. Second, we develop 90 real estate-specific question templates, addressing factual, inferential, and comparative queries. These templates are populated with real data to generate accurate QA pairs, with answers provided in Markdown, SQL-style, and natural language formats. Third, we annotate intent and slot labels by reverse engineering, classifying queries into 16 intent categories and identifying six slot categories based on entity types found in the query templates and table headers. Finally, we use in-context learning to rewrite the template-generated questions, enhancing their naturalness and diversity. After filtering, RETQA retains 20,762 QA pairs.

In summary, RETQA accurately mirrors the real-world challenges users encounter when querying real estate-related data, including open-domain queries, long-table queries, and multi-domain and multi-table queries, making RETQA a particularly demanding task. To address these challenges, we propose SLUTQA as a benchmark for RETQA, which improves the performance of large language models (LLMs) on TQA tasks by leveraging SLU labels through in-context learning (ICL).

The SLUTQA framework consists of three key modules: the SLU module, the SLU label-based Retrieval (SR) module, and the SLU label-based Filtering-Answer (SFA) module. For each query, the SLU module generates intent and slot labels at first. Depending on the availability of SLU labels, we use two approaches: (1) fine-tuning a BERT (Devlin et al. 2019) model when labels are sufficient, and (2) employing in-context learning with large language models (LLMs) when labels are limited (few-shot). The SR module then uses those SLU labels to create a query summary via in-context learning, improving retrieval accuracy by replacing the original query when searching for the relevant table caption by BM25 (Robertson and Walker 1994). Finally, the SFA module generates accurate SQL statements or refined Markdown answers based on the predicted SLU labels and retrieved tables. This framework leverages the strengths of large language models while integrating SLU task labels,

enhancing the performance of open-domain long-table question answering tasks. To the best of our knowledge, this is the first work to combine SLU tasks with TQA tasks.

Extensive experiments on RETQA demonstrate that SLUTQA significantly enhances LLM performance without requiring fine-tuning. Our dataset and framework offer valuable resources for advancing TQA research in the real estate domain, promoting further development in this field. The code and data will be released after the blind review. The main contributions of this paper are as follows:

- We introduce the first large-scale open-domain Chinese TQA dataset in the real estate domain, enriched with SLU labels. This dataset is challenging due to its long-table, multi-table, and open-domain characteristics.
- We propose the SLUTQA framework, which integrates LLMs with SLU tasks, significantly improving TQA accuracy on the RETQA dataset. To the best of our knowledge, this is the first work to combine SLU tasks with table QA tasks.
- Extensive experimental results show that SLUTQA improves the performance of existing LLMs on RETQA without the need for fine-tuning.

Related Works

Our work is related to two areas, i.e., TQA datasets, and TQA methods.

TQA Datasets

Early research, such as WikiTableQuestions (Pasupat and Liang 2015) and WikiSQL (Zhong, Xiong, and Socher 2017), collected tabular data from web sources like Wikipedia. However, WikiTableQuestions only provides text answers, and WikiSQL’s question descriptions are too vague to reliably locate relevant tables. The Spider dataset (Yu et al. 2018), designed for complex, cross-domain Text-to-SQL tasks, is not suited for open-domain scenarios and only offers SQL-format answers. NQ-TABLES (Herzig et al. 2021), the first open-domain table QA dataset, uses a reader model to extract answers from K candidate tables by selecting a single cell, limiting the need for complex

Dataset	Open Domain	# of QA pairs	# Tables	Answer format	Multi-table	SLU	Long Table
WikiSQL	✗	80654	24241	SQL	✗	✗	✗
WikiTableQuestion	✗	22033	2108	Text	✗	✗	✗
Spider	✗	10181	1020	SQL	✓	✗	✓
Open-WikiTable	✓	67023	24680	Text,SQL	✗	✗	✗
NQ-TABLES	✓	11628	169898	Text	✗	✗	-
RETQA (Ours)	✓	20762	4932	Text,SQL	✓	✓	✓

Table 1: Dataset Comparison, where “-” indicates data currently inaccessible. We define a dataset as a long table dataset if the average number of rows in the query-related tables exceeds 100.

reasoning (Kweon et al. 2023). Open-WikiTable (Kweon et al. 2023), based on WikiSQL and WikiTableQuestions, offers both SQL and text answers but typically features short tables. To the best of our knowledge, we are the first TQA dataset providing the SLU labels. The comparison of RETQA to other related datasets is shown in Table 1.

TQA Methods

Before the emergence of large language models (LLMs), researchers explored methods to combine tabular data with neural networks for natural language processing and data management tasks (Badaro, Saeed, and Papotti 2023; Fang et al. 2024; Shwartz-Ziv and Armon 2022). LLMs gained attention for their strong cross-task generalization, adapting to new tasks with minimal examples. (Chen 2023) first demonstrated LLMs’ ability to reason over tables through in-context learning. TAP4LLM (Sui et al. 2023) addresses noisy table data by enhancing sub-tables. OPENTAB retrieves tables, generates SQL as intermediate steps, and relies on a reader for final answers, though it may produce incorrect SQL (Kong et al. 2024). StructGPT (Jiang et al. 2023) uses interfaces to handle large structured data efficiently. MultiTabQA (Pal et al. 2023) introduces an LLM-based framework for answering questions across multiple tables, but it is limited to queries in closed domains. In comparison, our novel integration of LLMs with SLU tasks significantly enhances open-domain TQA accuracy.

Dataset Construction and Analysis

This section elaborates on how RETQA is created, including details on table collection, QA pair generation, intent and slot annotation, and query rewriting and quality control.

Table Collection

Our table data comes from publicly available real-world sources, including property data, real estate company financial data, and land auction data. The specific data sources and descriptions are as follows:

Property Data: We collected commercial housing transaction data from eight major Chinese cities (Beijing, Shanghai, Guangzhou, Shenzhen, Suzhou, Hangzhou, Nanjing, Wuhan) for 2019-2022, sourced from <http://www.fangdi.com.cn/>. The tables, keyed by development name, include district, average transaction prices, year and month of transaction dates, number of transactions,

developers, etc. This resulted in 4,825 standardized tables, captioned with the administrative region and year-month (e.g., “Residential Property Sales Price Table for Jiangning District of Nanjing City, May 2022”).

Real Estate Company Finance Data: Financial disclosures from Chinese real estate companies for 2019-2022 are collected from <https://aur.elmleaf.com.cn>. Each table, keyed by company name, includes total operating revenue, operating profit, total operating costs, total assets, total liabilities, state-owned status, credit bond status, and risk level. The data is organized by year, resulting in four tables captioned with the year (e.g., “China’s Real Estate Companies’ Financial Statements for 2020”).

Land Auction Data: Land auction information from 2016-2022 for the same eight cities is collected from city-specific public websites: Beijing¹, Shanghai², Guangzhou³, Shenzhen⁴, Suzhou⁵, Hangzhou⁶, Nanjing⁷, and Wuhan⁸. Tables, keyed by land parcel name, include details like development name, transaction date, affiliated group, total transaction price, floor area ratio, building density (%), green coverage ratio (%), etc. This yielded 103 tables, captioned by district (e.g., “Land Auction Information Table for Xuhui District, Shanghai”).

In total, our dataset integrates 4,932 tables across these three domains, with the longest table containing 465 rows. Each table is annotated with a standardized caption for retrieval. For more details about the table, see Appendix C.

QA Pair Generation

We utilize a template-based approach to generate QA pairs automatically. Specifically, we define 90 query templates covering various types of inquiries, including factual, inferential, and comparative questions, derived from extracted tables and create corresponding SQL templates for each query. Among these, 23 templates are designed to generate queries that require multi-table support. Given the complex

¹<https://zjw.beijing.gov.cn/>

²<https://biz.ghzyj.sh.gov.cn/shtdsc/wz>

³<https://zfcj.gz.gov.cn/zfcj/fyxx/fdctxmxx>

⁴<http://zjj.sz.gov.cn:8004/>

⁵<http://112.80.51.227:6600/index>

⁶<https://www.zjzrzyjy.com/portalBrowse/home>

⁷<http://www.landnj.cn>

⁸<http://www.whtdsc.com>

and data-intensive nature of the real estate market, we focus on generating queries that target longer tables to ensure that the dataset accurately reflects this complexity. Consequently, the average number of rows per query-related table is 252.9. Additional details on the templates are provided in Appendix A.

Using these templates, we generated 300 natural language queries and corresponding SQL statements per template, resulting in a total of 27,000 pairs. We then filtered out QA pairs with duplicate queries, non-executable SQL statements, and SQL statements that return empty results, leaving us with 20,471 QA pairs. For each remaining pair, we also provide a matching Markdown-formatted input table (with multiple tables concatenated into a single row) and the corresponding output answer table. Since the real estate data is collected from Chinese public websites, all queries are generated in Chinese. We plan to publish an English version soon.

Intent and Slot Annotation To facilitate researchers effectively extracting key information from queries and achieve accurate retrieval and answers, RETQA includes additional Spoken Language Understanding (SLU) labels. SLU consists of two types of labels: intent and slot labels, commonly used in task-oriented dialogue systems. Intent labels identify the user’s query intent, while slot labels extract key information pertinent to that intent.

We annotate the intent and slot labels through reverse engineering. Specifically, for intent labels, we categorized the queries into 16 types based on the 90 templates mentioned earlier. Some queries in the templates may involve multiple intents. For example, in the query, “Which of the top 5 best-selling residential compounds in Bao’an District, Shenzhen, has the lowest housing price”, both average price and sales volume information are required. In this case, the query would be labeled with two intents: “real estate project sales volume query” and “real estate project average transaction price query.” For more details and statistics about the intent labels, please see Appendix C.

For slot labels, we categorize the entities into six types: “city”, “district”, “development name”, “company name”, “year”, and “month”, based on the entity types in the query templates and corresponding table headers. Following previous work on SLU tasks (Qin et al. 2022), we adopt the Inside–Outside–Beginning (IOB) tagging format (Ramshaw and Marcus 1999). In this format, the B- prefix indicates the beginning of a slot chunk, the I- prefix indicates that the tag is inside a slot chunk, and the O tag indicates that a token does not belong to any slot chunk. An example of SLU labels is shown in Figure 1.

LLM-based Query Rewriting and Quality Control Despite our extensive template library, the sentences generated from these templates can be grammatically monotonous, differing significantly from real human queries. This disparity may cause models trained on template-based data to struggle with actual user queries. To address this, we use large language models (LLMs) to rewrite the queries generated by templates, creating more diverse and human-like expressions while retaining their original meaning. Tech-

niques such as synonym replacement and sentence inversion are employed in this process.

We used Qwen2 72B (Yang et al. 2024) for rewriting and DeepSeek-V1 (Dense-67B) (DeepSeek-AI 2024) to evaluate the results. The evaluation scores questions on a scale from 0 (“completely like template generation”) to 5 (“like human writing”). After rewriting, the average score increased from 2.56 to 2.95, indicating a significant improvement in naturalness and diversity. The distribution of scores before and after rewriting is shown in Appendix B.

After rewriting, we manually reviewed and removed 291 incorrect QA pairs, leaving 20,762 pairs in the RETQA dataset, and split them into training, testing, and validation sets at a ratio of 0.8:0.1:0.1. To ensure that each template was represented in each set, we perform sampling based on the templates. Finally, for each query, RETQA provides the corresponding intent labels, slot labels, and the caption of the relevant table(s), as well as answers in three formats: Markdown, SQL-style (with both the SQL statements and their resulting answers), and natural language. Detailed statistics of RETQA are shown in Appendix C.

Method

In this section, we introduce the SLUTQA framework, which consists of three key modules: the SLU module, the SLU label-based Retrieval (SR) module, and the SLU label-based Filtering-Answer (SFA) module. SLUTQA aims to enhance the performance of LLMs on TQA tasks through SLU labels-based in-context learning (ICL). The general framework of SLUTQA is shown in Figure 2.

SLU Module

SLU involves two label types, i.e., intent and slot, commonly used in task-oriented dialogue systems. Intent labels identify the user’s query intent, while slot labels extract relevant key information. In this paper, we design a SLU module to predict the SLU labels of each query, and leverage these labels to enhance information extraction and ensure accurate retrieval and answers within the SLUQA framework.

Based on the availability of labeled SLU data, we design two approaches to implement the SLU module: (1) fine-tuning a BERT model (Devlin et al. 2019) when sufficient labeled data are available, and (2) utilizing ICL with LLMs in scenarios where labeled data are limited (few-shot).

In the first scenario, given the current query $\mathbf{X} = \{x_1, \dots, x_n\}$ as input, we utilize a pre-trained BERT⁹ as the encoder. The model returns the hidden states $H = \{h_{cls}, h_1, \dots, h_n \in \mathbb{R}^{d_{model}}\}$, where CLS is a special token representing the entire input sequence within BERT, and d_{model} is the output dimension of BERT.

Then, we predict the intent labels, number of intents, and slot labels by:

$$\begin{aligned} y^I &= \text{Sigmoid}(\mathbf{W}^I \cdot \mathbf{h}_{cls} + \mathbf{b}^I), \\ y^N &= \text{Softmax}(\mathbf{W}^N \cdot \mathbf{h}_{cls} + \mathbf{b}^N), \\ y_j^S &= \text{Softmax}(\mathbf{W}^S \cdot \mathbf{h}_j + \mathbf{b}^S), \end{aligned} \quad (1)$$

⁹<https://huggingface.co/google-bert/bert-base-uncased>

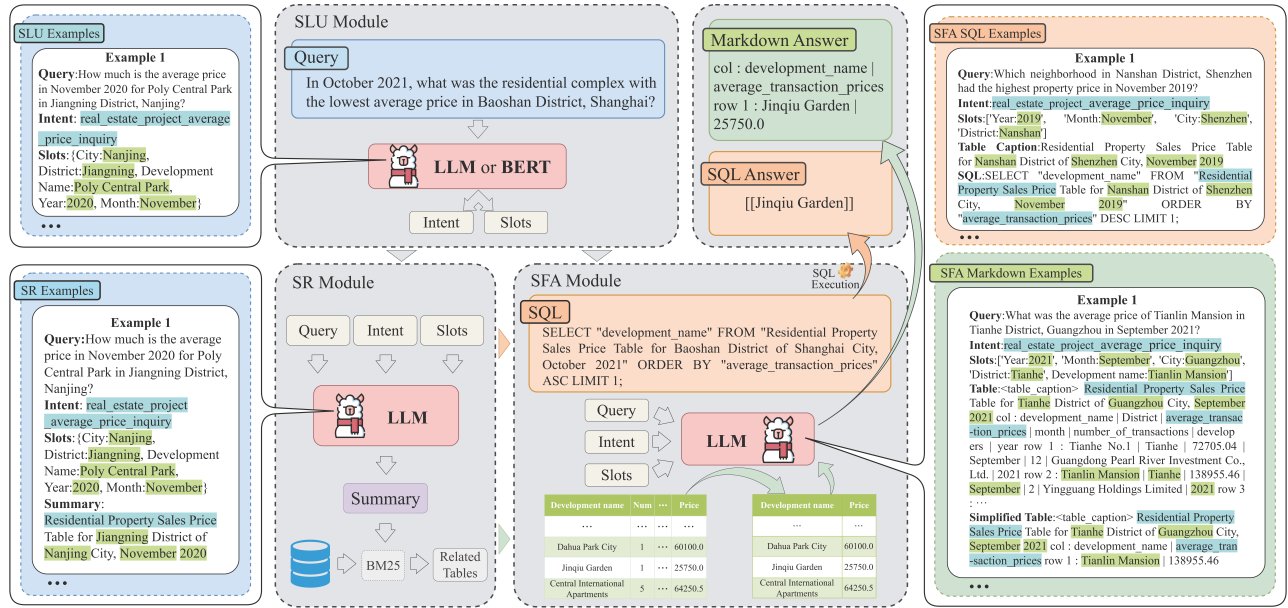


Figure 2: General framework of SLUTQA.

where $\mathbf{y}^I \in \mathbb{R}^{d_i}$, $\mathbf{y}^N \in \mathbb{R}^{d_n}$, and $\mathbf{y}^S = \mathbf{y}_1^S, \dots, \mathbf{y}_n^S \in \mathbb{R}^{d_s}$ represent the predicted results of intents, number of intent, and slots, respectively. $\mathbf{W}^I \in \mathbb{R}^{d_i \times d_{model}}$, $\mathbf{W}^N \in \mathbb{R}^{d_n \times d_{model}}$, and $\mathbf{W}^S \in \mathbb{R}^{d_s \times d_{model}}$ are fully connected matrices, $\mathbf{b}^I \in \mathbb{R}^{d_i}$, $\mathbf{b}^N \in \mathbb{R}^{d_n}$ and $\mathbf{b}^S \in \mathbb{R}^{d_s}$ are bias vectors. In this paper, $d_i = 16$, $d_n = 2$, and $d_s = 13$ (6 kinds of slots with B-tags, I-tags, and the O-tag), representing the categories of intent labels, number of intents, and categories of slot labels, respectively.

During fine-tuning, cross-entropy loss is employed for both intent number prediction and slot filling, while binary cross-entropy is used for intent detection. During inference, the top 1 or 2 predicted intents are selected based on the predicted number of intents.

In the second scenario, where labeled SLU data are limited (few-shot) and insufficient for training or fine-tuning a model, we predict SLU labels using ICL in LLMs. Specifically, we randomly select 22 examples from the training set to serve as prompt examples for the LLMs. As illustrated in Figure 2 (top left, SLU Examples), the LLM leverages these few-shot examples as prompts to jointly predict intent and slot labels for each input query.

SR Module

Although our dataset provides corresponding tables (i.e., gold tables) for each query, in real-world open-domain scenarios, these tables must be retrieved based on the query, and the retrieval accuracy directly affects overall TQA performance. Previous works (Kong et al. 2024) often use BM25 (Robertson and Zaragoza 2009) for table retrieval due to its scalability and competitive performance. However, in the real estate data scenario, user queries can be diverse and non-standard, while table captions are often similar (e.g., differing only by years or months). This means that relying solely

on BM25 for table retrieval may not capture the critical information needed, leading to lower accuracy in our scenario.

To address this issue, we propose the SR module, which utilizes ICL with SLU labels to enable the large LLMs to generate a query summary in the form of a table caption. For each query, five examples from the training set that share the same intent (as predicted by the SLU module) are selected and used as prompts for the LLM. As illustrated in Figure 2 (bottom left, SR Examples), each example comprises four components: the query, intent(s), slots, and query summary, with the query summary corresponding to the table caption(s) in the example, which may include multiple tables. This approach enables the LLM to learn to generate a summary in the form of a table caption relevant to the query. The intent and slot information allow the LLM to more accurately capture the key details in the query, resulting in a more precise summary. Finally, the SR-generated summary, rather than the original query, is used to retrieve the table caption via BM25. It is important to note that the generated summary may consist of multiple entries, with each summary being used to retrieve one table using BM25, selecting the top result for each.

SFA Module

After retrieving the corresponding tables from the SR module, we designed the SFA module to generate final answers in two formats: SQL-style and markdown.

In markdown formats, existing methods typically flatten entire tables into a single row for processing by large language models (LLMs). However, RETQA includes many lengthy tables, and current LLMs, with their limited input lengths (usually 32k tokens), struggle to effectively handle queries involving these extensive tables. To address this issue, the SFA module for markdown formats is designed to

Model	Method	Table EM(%)	Row EM(%)			Column EM(%)			Cell EM(%)		
			P	R	F1	P	R	F1	P	R	F1
Qwen2 7b	Vanilla	15.05	24.43	31.71	27.60	25.10	22.11	23.51	36.19	43.35	39.45
	SLUTQA (ICL)	16.65	26.73	25.34	26.02	23.36	26.06	24.64	40.67	40.45	40.56
	SLUTQA (FT)	18.94	32.36	29.61	30.92	23.14	28.90	25.70	45.07	43.24	44.14
GLM4 9b	Vanilla	4.73	26.55	45.71	33.59	31.17	30.36	30.76	35.77	59.42	44.66
	SLUTQA (ICL)	9.15	34.88	40.42	37.44	28.19	34.25	30.92	44.66	55.80	49.61
	SLUTQA (FT)	9.08	37.62	43.44	40.32	28.73	33.93	31.11	47.75	58.01	52.38
Qwen2 72b	Vanilla	11.14	36.27	54.54	43.57	32.16	32.93	32.54	44.64	58.46	50.62
	SLUTQA (ICL)	14.81	49.79	53.74	51.69	44.01	42.75	43.37	64.51	67.68	66.06
	SLUTQA (FT)	15.02	51.26	55.61	53.34	43.39	41.90	42.63	65.01	68.29	66.61

Table 2: Overall performance of markdown format answer.

Model	Method	ECR(%)	pass@1
Qwen2 7b	Vanilla	60.11	39.55
	SLUTQA (ICL)	85.46	64.46
	SLUTQA (FT)	87.28	70.16
Qwen2 72b	Vanilla	68.49	51.40
	SLUTQA (ICL)	92.06	75.65
	SLUTQA (FT)	97.23	82.71
GLM4 9b	Vanilla	56.62	34.52
	SLUTQA (ICL)	84.97	60.99
	SLUTQA (FT)	89.90	66.81

Table 3: Overall performance of SQL format answer.

filter out rows and columns that are irrelevant to the query. This is achieved with the help of SLU labels, which contain key information that guides the LLM to focus on the relevant portions of the table. Specifically, the SFA module operates in two steps. First, the original markdown table is simplified by removing irrelevant rows and columns based on the identified slots and intent. Second, the simplified table is used to generate the final output. This process is facilitated through ICL. For each query, 2-5 examples (according to the total length) from the training set, all sharing the same intent, are selected. As illustrated in Figure 2 (bottom right, SFA Markdown Examples), the first stage of the process includes examples consisting of five components: the query, intent(s), slots, original table, and simplified table. In the second stage, the examples include the query, intent(s), slots, simplified table, and the final answer.

For SQL-style answers, the primary challenge is accurately extracting key information from queries to generate the corresponding SQL statements. To address this, the SLA module for SQL formats leverages SLU labels and ICL to generate SQL statements. As with the markdown process, we select five examples from the training set for each query, all sharing the same intent. As illustrated in Figure 2 (top right, SFA SQL Examples), these examples consist of the query, intent(s), slots, retrieved table caption(s), and SQL statements. The final output is produced by executing the generated SQL statement.

Experiments

In this section, we first introduce the experiment setup. Then, we show the experiment results and conduct ablation

studies.

Experimental Settings and Baselines

Evaluation Metrics: In this paper, we provide answers to each query in three formats: Markdown, SQL-style, and natural language. However, since natural language output is subjective, we focus on evaluating the two objective formats, Markdown and SQL-style.

For SQL-style answers, we assess performance using Executable Code Ratio (ECR) and Pass Rate (pass@1), as outlined in (He et al. 2024). For Markdown answers, we evaluate using Table Exact Match accuracy (Table EM) and Precision (P), Recall (R), and F1 score for exact matches of rows, columns, and cells, following the metrics described in (Pal et al. 2023). Additionally, for table retrieval performance, we also employ Precision, Recall, and F1 score as evaluation metrics.

Baselines: To prove that our designed SLUTQA could enhance the performance of existing LLMs on the RETQA dataset, we utilize our framework on three GPT family models, i.e., Qwen2-7b, Qwen2-72b (Yang et al. 2024), and Glm4-9b (GLM et al. 2024).

Result and Analysis

Main Results: To evaluate the performance of SLUTQA on RETQA, we generated markdown and SQL-style answers using baseline LLMs through two approaches: the vanilla implementation and our SLUTQA framework. In the vanilla implementation, we generated a query summary for BM25 retrieval and then produced the final answers using ICL with five randomly selected examples. The key difference is that the vanilla approach does not include SLU labels and lacks a simplification step for markdown formatting.

The overall performance of the markdown and SQL-style answers is presented in Tables 2 and 3, respectively. In these tables, “Vanilla” refers to the baseline implementation, “SLUTQA (ICL)” denotes the SLUTQA framework with the SR module implemented via in-context learning (ICL), and “SLUTQA (FT)” represents the SLUTQA framework with the SR module implemented by fine-tuning a BERT model. The results clearly demonstrate that our SLUTQA framework outperforms all baseline methods in both markdown and SQL formats, and further confirms that incorporating SLU labels can significantly enhance the performance of LLMs in open-domain TQA tasks.

Method	P	R	F1	
BM25	Top 1	76.86	56.02	64.80
Qwen2 7B	Vanilla	93.87	89.09	91.42
	SR(ICL)	95.80	92.32	94.03
	SR(FT)	97.53	95.05	96.27
Qwen2 72B	Vanilla	97.23	96.30	96.76
	SR(ICL)	97.66	97.35	97.50
	SR(FT)	97.85	97.57	97.71
GLM4 9B	Vanilla	93.73	92.02	92.87
	SR(ICL)	95.12	92.75	93.92
	SR(FT)	97.67	94.29	95.95

Table 4: Table retrieval performance of different methods.

Model	Method	Cell EM(%)		
		P	R	F1
Qwen2 7b	Vanilla	37.67	43.58	40.41
	Simplified	34.07	37.88	35.87
	SFA	39.62	44.02	41.71
GLM4 9b	Vanilla	39.85	61.96	48.51
	Simplified	42.25	57.75	48.80
	SFA	51.89	62.87	56.85
Qwen2 72b	Vanilla	45.72	58.97	51.51
	Simplified	62.35	69.20	65.59
	SFA	66.36	70.27	68.26

Table 5: Ablation study of SFA module in markdown format.

Specifically, for markdown format answers, comparing SLUTQA (FT) with the vanilla implementation, Table EM accuracy improves by 3.89%, 4.35%, and 3.88% for Qwen2 7b, GLM4 9b, and Qwen2 72b, respectively. Even when comparing SLUTQA (ICL), which does not fine-tune with a large number of SLU labels, with the vanilla implementation, Table EM accuracy still improves by 1.6%, 4.42%, and 3.67% for Qwen2 7b, GLM4 9b, and Qwen2 72b, respectively. For SQL format answers, both SLUTQA (FT) and SLUTQA (ICL) outperform the vanilla implementation by more than 20% on both ECR and pass@1 scores.

These results indicate that: (1) Current LLMs are capable of generating accurate SLU labels with just few-shot (22 in practice) examples. (2) SLU labels facilitate more precise parsing of user intent and relevant information, significantly enhancing the overall performance of open-domain TQA tasks. Detailed results of the SLU tasks are provided in Appendix D. This also demonstrates that incorporating traditional SLU labels into open-domain TQA tasks is a valid and effective approach, offering valuable insights for future TQA research.

Ablation Study: We conducted an ablation study to evaluate the effectiveness of the SR and SFA modules in SLUTQA.

First, we evaluated the SR module’s impact on retrieval performance, as shown in Table 4. We compared several methods: BM25, which retrieves the top 1 table based on the original query; the vanilla approach, which generates query summaries replacing the original query for BM25 retrieval; and our SR module. In SR (ICL) and SR (FT), summaries are generated using our SR module, with SLU labels pre-

Model	Method	ECR(%)	pass@1
Qwen2 7b	Vanilla	64.99	42.24
	SFA	90.17	74.58
Qwen2 72b	Vanilla	71.76	53.70
	SFA	98.86	86.26
GLM4 9b	Vanilla	60.01	35.25
	SFA	90.36	69.65

Table 6: Ablation study of SFA module in SQL format.

dicted by ICL or a fine-tuned (FT) BERT model. The results show that BM25 performs significantly worse than the other methods, highlighting the effectiveness of LLM-generated table summaries for TQA tasks. Additionally, both SR (ICL) and SR (FT) outperform the vanilla approach across all LLMs, demonstrating that SLU labels enhance retrieval accuracy.

Next, we evaluated the SFA module. To isolate its impact, we used the gold table(s) and ground-truth SLU labels, ensuring the SFA module’s performance is assessed independently of retrieval accuracy and SLU prediction.

The markdown results are shown in Table 5. Due to the space limit, we only show the performance of exact matches of cells. Here “Vanilla” refers to the vanilla implementation described in the Main Results section, where the LLM generates the final answer from the query and gold markdown table(s) using in-context learning with 5 randomly selected examples. “Simplified” is an ICL-based method similar to the SFA module but without SLU labels in the prompts. “SFA” represents our SFA module. The results demonstrate that the SFA module outperforms both Vanilla and Simplified across all LLMs, indicating that SLU labels help simplify tables and reduce noise, addressing the challenges of long tables. Notably, Vanilla outperforms Simplified on Qwen2 7B, suggesting that without SLU labels, direct table simplification through ICL can inadvertently remove important information, reducing performance.

The SQL results are shown in Table 6. Again, “Vanilla” refers to the vanilla implementation in the Main Results section, while “SFA” denotes our SFA module. The results show that the SFA module significantly outperforms Vanilla across all LLMs, indicating that SLU labels help LLMs capture key information and generate more accurate SQL statements.

Conclusion

In this paper, we introduce RETQA, the first large-scale open-domain Chinese TQA dataset for the real estate sector, comprising 4,932 tables and 20,762 QA pairs. RETQA addresses the lack of specialized datasets by providing a challenging resource for TQA research, particularly in handling open-domain, long tables and multi-domain queries. To enhance performance, we developed the SLUTQA framework, which leverages LLMs and SLU labels for improved retrieval and answer generation. Extensive experiments show that SLUTQA significantly boosts LLM performance, offering valuable insights and resources for future TQA research.

Acknowledgements

This work was supported in part by the Chinese National Research Fund (NSFC) under Grant 62272050 and Grant 62302048; Zhuhai Science-Tech Innovation Bureau under Grant No. 2320004002772, and in part by the Interdisciplinary Intelligence SuperComputer Center of Beijing Normal University (Zhuhai). The acquisition of real estate data is supported by Elmleaf Limited (Shanghai).

References

- Badaro, G.; Saeed, M.; and Papotti, P. 2023. Transformers for Tabular Data Representation: A Survey of Models and Applications. *Transactions of the Association for Computational Linguistics*, 11: 227–249.
- Chen, W. 2023. Large Language Models are few (1)-shot Table Reasoners. In *Findings of the Association for Computational Linguistics: EAACL 2023*, 1120–1130.
- Cheng, L.; Yang, W.; and Jia, W. 2023. A scope sensitive and result attentive model for multi-intent spoken language understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 12691–12699. AAAI Press.
- DeepSeek-AI. 2024. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism. *arXiv preprint arXiv:2401.02954*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Fang, X.; Xu, W.; Tan, F. A.; Zhang, J.; Hu, Z.; Qi, Y. J.; Nickleach, S.; Socolinsky, D.; Sengamedu, S.; Faloutsos, C.; et al. 2024. Large language models (LLMs) on tabular data: Prediction, generation, and understanding—a survey. *arXiv preprint arXiv:2402.17944*.
- GLM, T.; Zeng, A.; Xu, B.; Wang, B.; Zhang, C.; Yin, D.; Rojas, D.; Feng, G.; Zhao, H.; Lai, H.; et al. 2024. ChatGLM: A Family of Large Language Models from GLM-130B to GLM-4 All Tools. *arXiv preprint arXiv:2406.12793*.
- He, X.; Zhou, M.; Xu, X.; Ma, X.; Ding, R.; Du, L.; Gao, Y.; Jia, R.; Chen, X.; Han, S.; et al. 2024. Text2analysis: A benchmark of table question answering with advanced data analysis and unclear queries. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 18206–18215. AAAI Press.
- Herzig, J.; Mueller, T.; Krichene, S.; and Eisenschlos, J. 2021. Open Domain Question Answering over Tables via Dense Retrieval. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 512–519.
- Jiang, J.; Zhou, K.; Dong, Z.; Ye, K.; Zhao, W. X.; and Wen, J.-R. 2023. StructGPT: A General Framework for Large Language Model to Reason over Structured Data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9237–9251.
- Kong, K.; Zhang, J.; Shen, Z.; Srinivasan, B.; Lei, C.; Faloutsos, C.; Rangwala, H.; and Karypis, G. 2024. OpenTab: Advancing Large Language Models as Open-domain Table Reasoners. *CoRR*, abs/2402.14361.
- Kweon, S.; Kwon, Y.; Cho, S.; Jo, Y.; and Choi, E. 2023. Open-WikiTable: Dataset for Open Domain Question Answering with Complex Reasoning over Table. In *Findings of the Association for Computational Linguistics: ACL 2023*, 8285–8297.
- Pal, V.; Yates, A.; Kanoulas, E.; and de Rijke, M. 2023. MultiTabQA: Generating Tabular Answers for Multi-Table Question Answering. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 6322–6334.
- Pasupat, P.; and Liang, P. 2015. Compositional Semantic Parsing on Semi-Structured Tables. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1470–1480.
- Qin, L.; Wei, F.; Chen, Q.; Zhou, J.; Huang, S.; Si, J.; Lu, W.; and Che, W. 2024. CroPrompt: Cross-task Interactive Prompting for Zero-shot Spoken Language Understanding. *arXiv preprint arXiv:2406.10505*.
- Qin, L.; Wei, F.; Ni, M.; Zhang, Y.; Che, W.; Li, Y.; and Liu, T. 2022. Multi-domain spoken language understanding using domain-and task-aware parameterization. *Transactions on Asian and Low-Resource Language Information Processing*, 21(4): 1–17.
- Qin, L.; Xie, T.; Che, W.; and Liu, T. 2021. A Survey on Spoken Language Understanding: Recent Advances and New Frontiers. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 4577–4584. ijcai.org.
- Ramshaw, L. A.; and Marcus, M. P. 1999. Text chunking using transformation-based learning. In *Natural language processing using very large corpora*, 157–176. Springer.
- Robertson, S.; and Zaragoza, H. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends® in Information Retrieval*, 3(4): 333–389.
- Robertson, S. E.; and Walker, S. 1994. Some simple effective approximations to the 2-poisson model for probabilistic weighted retrieval. In *SIGIR'94: Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval, organised by Dublin City University*, 232–241. Springer.
- Shwartz-Ziv, R.; and Armon, A. 2022. Tabular data: Deep learning is not all you need. *Information Fusion*, 81: 84–90.
- Sui, Y.; Zou, J.; Zhou, M.; He, X.; Du, L.; Han, S.; and Zhang, D. 2023. Tap4llm: Table provider on sampling, augmenting, and packing semi-structured data for large language model reasoning. *arXiv preprint arXiv:2312.09039*.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Yu, T.; Zhang, R.; Yang, K.; Yasunaga, M.; Wang, D.; Li, Z.; Ma, J.; Li, I.; Yao, Q.; Roman, S.; et al. 2018. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3911–3921.

Zhong, V.; Xiong, C.; and Socher, R. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.

Zhu, Z.; Cheng, X.; An, H.; Wang, Z.; Chen, D.; and Huang, Z. 2024. Zero-Shot Spoken Language Understanding via Large Language Models: A Preliminary Study. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, 17877–17883.