

MURRE: Multi-Hop Table Retrieval with Removal for Open-Domain Text-to-SQL

Xuanliang Zhang¹, Dingzirui Wang¹, Longxu Dou², Qingfu Zhu¹, Wanxiang Che^{1*}

¹Harbin Institute of Technology

²Individual Researcher

{xuanliangzhang, dzrwang, qfzhu, car}@ir.hit.edu.cn

doulongxu@gmail.com

Abstract

The open-domain text-to-SQL task aims to retrieve question-relevant tables from massive databases and generate SQL. However, the performance of current methods is constrained by single-hop retrieval, and existing multi-hop retrieval of open-domain question answering is not directly applicable due to the tendency to retrieve tables similar to the retrieved ones but irrelevant to the question. Since the questions in text-to-SQL usually contain all required information, while previous multi-hop retrieval supplements the questions with retrieved documents. Therefore, we propose the multi-hop table retrieval with removal (MURRE), which removes previously retrieved information from the question to guide the retriever towards unretrieved relevant tables. Our experiments on two open-domain text-to-SQL datasets demonstrate an average improvement of 5.7% over the previous state-of-the-art results.¹

1 Introduction

Text-to-SQL, which simplifies database access and facilitates efficient data querying, is an important task in natural language processing (Qin et al., 2022). Unlike the previous text-to-SQL task that provides the question-relevant database tables² (Yu et al., 2018; Shi et al., 2024), a more realistic scenario involves open-domain text-to-SQL, where user questions are transformed into SQL across vast databases. Therefore, open-domain text-to-SQL necessitates two main steps: retrieving relevant tables and generating SQL based on the retrieved tables and user question (Kothiyari et al., 2023).

To bridge the semantic gap between natural language user questions and structured database tables, CRUSH (Kothiyari et al., 2023) rewrites the user

*Corresponding Author

¹Our code and data are available at <https://github.com/zhxlia/Murre>.

²For brevity, we refer to database tables relevant to the question as **relevant tables** in this paper.

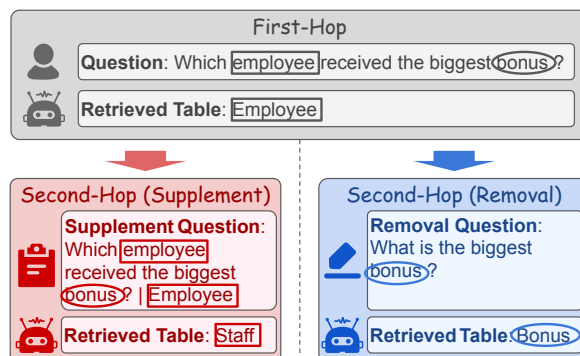


Figure 1: Comparison of multi-hop retrieval with supplementary and removal. The same shapes denote similar tables. Multi-hop retrieval with supplementary retrieves tables similar to those already retrieved, even if they are irrelevant to the question. In contrast, we employ multi-hop retrieval with removal which can successfully retrieve other relevant tables in the second hop.

question into potentially relevant tables for retrieval. However, CRUSH employs single-hop retrieval, limiting its performance, as the retrieval of certain tables may depend on others, similar to multi-hop retrieval in open-domain question answering (QA) (Feldman and El-Yaniv, 2019; Xiong et al., 2021). Nevertheless, directly adapting previous multi-hop retrieval methods could easily *retrieve tables that are similar to those retrieved in previous hops but irrelevant to the question* since previous multi-hop retrieval usually supplements the retrieved documents into the user question (Shao et al., 2023; Chen et al., 2024a; Jeong et al., 2024), as illustrated in Figure 1. However, open-domain text-to-SQL questions typically contain all necessary information (Liu et al., 2021; Dou et al., 2023). For instance, in the left part of Figure 1, the irrelevant table "staff" is retrieved because it is similar to the "employee" retrieved in the first hop.

To solve the above problem, we focus on exploring the multi-hop table retrieval method applicable to the open-domain text-to-SQL task, from the following two aspects: (i) We analyze why the multi-hop retrieval methods of open-domain QA cannot

be directly applied to open-domain text-to-SQL. (ii) We propose a multi-hop table retrieval method based on removal for open-domain text-to-SQL, to ensure that the retrieved tables of each hop are relevant to the user question rather than just similar to the previously retrieved tables.

We first present that multi-hop retrieval of open-domain QA cannot be directly used for open-domain text-to-SQL. We discuss that existing multi-hop retrieval methods are mainly to supplement the retrieved documents to the user question. However, open-domain text-to-SQL needs to remove retrieved information rather than supplement it because all the necessary information is involved in the user question. We conduct analysis experiments to prove our point of view.

Based on the above discussion, we propose a method called **MULTI**-hop table **RE**trieval with **RE**moval (MURRE), to enhance the performance of open-domain text-to-SQL through multi-hop retrieval. As presented in the right part of Figure 1, MURRE searches relevant tables employing multi-hop retrieval, by removing the retrieved table information from the question to ensure retrieved tables relevant to the user question.

To validate the effectiveness of MURRE, we conduct experiments on two datasets, SpiderUnion and BirdUnion, which are open-domain versions of text-to-SQL datasets Spider (Yu et al., 2018) and Bird (Li et al., 2023b). MURRE achieves an average improvement of 5.7% compared to the previous state-of-the-art (SOTA) results, demonstrating its effectiveness. Additionally, we calculate the average rank of relevant tables to demonstrate that MURRE can effectively retrieve relevant tables that are dissimilar to the previously retrieved tables, confirming that our method can indeed remove the retrieved information.

Our contributions are as follows:

- To enhance the retrieval performance in open-domain text-to-SQL, we discuss why multi-hop retrieval in open-domain QA can not be directly applied to text-to-SQL.
- To ensure that the retrieved tables are relevant to the user question, we present MURRE, prompting LLMs to remove the retrieved information from the question at each hop.
- To demonstrate the effectiveness of MURRE, we conduct experiments on the SpiderUnion and BirdUnion datasets, achieving an average improvement of 5.7% compared with the previous SOTA results, proving its effectiveness.

2 Analysis

In this section, we present that **the multi-hop retrieval of open-domain QA cannot be directly applied to the open-domain text-to-SQL**. We first discuss the multi-hop retrieval methods in open-domain QA. Based on the above discussion, we analyze and conduct experiments to show the inapplicability of existing multi-hop retrieval.

2.1 Multi-Hop Retrieval of Open-Domain QA

Multi-hop retrieval refers to retrieving with multiple hops, where each hop uses the information that is retrieved in previous hops (Xiong et al., 2021; Lee et al., 2022). Current multi-hop retrieval supplements the question with the retrieved information, guiding to identify the documents relevant to the currently retrieved documents (Chen et al., 2024a; Zhao et al., 2023a; Press et al., 2023; Li et al., 2023c; Jeong et al., 2024), since the user question could not contain all the information needed to retrieve all relevant documents. For example, for the question *"Ralph Hefferline was a psychology professor at a university that is located in what city?"*, we should first retrieve the university where he works, supplement it into the question, and then search for the city where the university is located.

2.2 Previous Multi-Hop Retrieval is Unsuitable for Open-Domain Text-to-SQL

Unlike open-domain QA, the questions in open-domain text-to-SQL do not need retrieved information to supplement because they usually contain all the required information. For instance, in Figure 1, the question *"Which employee received the biggest bonus?"* contains all the relevant table names. Therefore, we propose to remove the retrieved information from the question, to retrieve the unretrieved relevant tables. Table 4 presents that multi-hop retrieval with supplementary even underperforms single-hop retrieval, proving the necessity of removing retrieved information.

3 Methodology

3.1 Task Definition

Our work mainly focuses on the open-domain text-to-SQL task, which can be formally defined as: Given a user question q , database tables $T = \{t_i\}$ and a number of retrieved tables N , suppose the tables relevant to q are $T^q = \{t_i^q\}$, MURRE aims to retrieve N tables T_N^q , where $T^q \subseteq T_N^q$.

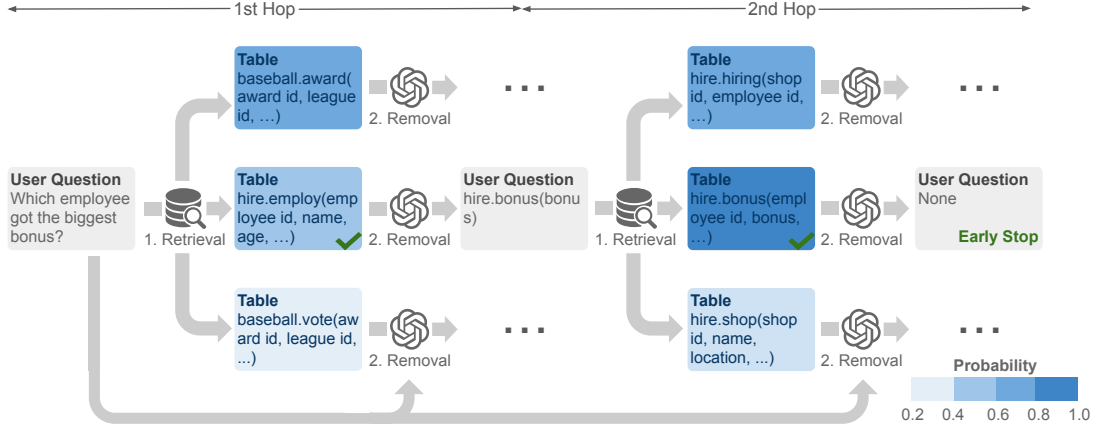


Figure 2: An overview of MURRE. Each hop consists of: (i) **Retrieval**: retrieving tables similar to the question; (ii) **Removal**: removing the retrieved information from the user question and representing the unretrieved information in the tabular format with LLM. We employ the beam search paradigm maintaining multiple retrievals at each hop. The color depth represents the probability that the table is relevant to the question of the hop, and ✓ denotes the relevant table. We demonstrate an example that early stops at the second hop for brevity.

3.2 Overview

The overview of MURRE is illustrated in Figure 2. To reduce the error cascades in multi-hop retrieval, we adopt the beam search paradigm inspired by Zhang et al. (2024). In MURRE, each hop can be divided into two phases: **Retrieval** (§3.3) and **Removal** (§3.4). In each hop, we first **retrieve** the relevant tables by calculating their probability that is relevant to the question. To avoid retrieving similar to the previous retrieved tables but irrelevant tables, we **remove** the retrieved information from the user question using LLM. MURRE repeats the Retrieval and Removal phases until reaching the maximum hop limit H or meeting the early stop condition. After the multi-hop retrieval, we **score** (§3.5) each table based on its probability that is relevant to the user question and select the top- N tables as the input for generating SQL. In Appendix A, we explore how to handle the situation when no table in the database is relevant to the user question.

3.3 Retrieval

The Retrieval phase aims to identify B tables corresponding to the user question $q^{h,b}$, where B is the beam size and $q^{h,b}$ denotes the user question at hop h and beam b . We embed the user question and each table into vectors and then compute their relevance probability to the user question. Let $\text{Emb}(x)$ represent the embedding vector of x , the probability is expressed as Equation 3.1.

$$\hat{P}(t_i|q^{h,b}) = \text{Norm}\left(\frac{\text{Emb}(t_i) \cdot \text{Emb}(q^{h,b})}{|\text{Emb}(t_i)| |\text{Emb}(q^{h,b})|}\right) \quad (3.1)$$

We compute the probability using the cosine similarity between the question and table vectors, applying the Norm function to ensure the conditional probabilities sum to one. The detailed representations of the table and the normalization method are provided in Appendix B and Appendix C, respectively. We then select the top B tables with the highest probabilities as the retrieval results for the current hop question $q^{h,b}$.

In the previous hop of retrieval, there are B beams and each beam corresponds to B retrieved tables obtained at the current hop, resulting in a total of $B \times B$ retrieval results. Following the beam search paradigm, we then choose B results from these for the subsequent Removal phase, with the selection method detailed in §3.5.

3.4 Removal

The Removal phase is designed to mitigate the retrieval of tables that are similar to previously retrieved ones yet irrelevant to the user question. To achieve the above goal, this phase prompts LLMs to remove the information of retrieved tables in $\text{Path}^{h,b}$ (which includes all retrieved tables from hop 1 to hop h on the path of beam b) from the user question q , and express the unretrieved information in the form of a table with some demonstrations, following the previous work (Kothiyari et al., 2023). The removal question is used to guide the retriever for the next hop.

Considering that different user questions require varying numbers of tables, to prevent additional hops from introducing errors, we instruct LLMs to

assess whether the retrieved tables in $\text{Path}^{h,b}$ are sufficient to answer q , i.e., early stop. In detail, we prompt the LLM to generate a special mark when removing the retrieved information (e.g., "None" in Figure 2) to indicate that the retrieved tables are sufficient to answer the user question q , ceasing further retrieval as soon as this early stop mark is produced. The prompts used for removal are presented in Appendix B.

3.5 Score

The aforementioned multi-hop retrieval process maintains B paths at each hop, containing tables that could exceed or fall short of the required number of tables N . Therefore, this phase aims to score all the tables based on their probabilities to retain the most relevant tables used for generating SQL, which consists of two parts: (i) scoring the retrieval path $\text{Score_Path}(\text{Path}^{h,b})$; (ii) scoring the retrieved table $\text{Score_Table}(t_i)$.

Retrieval Path Score First, we address the calculation of the retrieval path score, where each node on the path corresponds to the retrieved table on a beam of one hop. As discussed in §3.3, the score of each retrieval path represents the probability that the last table in the path is retrieved, given its corresponding question at the last hop. Following the derivation in Appendix D, the score of a retrieval path is computed as the product of all the probabilities \hat{P} in the path.

Table Score Building on the retrieval path score, we describe the calculation to score a retrieved table t_i . Since each table could have multiple scores across various hops and beam retrievals, we propose a table scoring algorithm to effectively integrate these scores. Considering the potential interrelation of question-relevant tables, we aim to ensure that all retrieved tables collectively are the most relevant to the question. Thus, the higher the retrieval path score, the higher the score of the tables in the path. Specifically, let Path_{t_i} denote all retrieval paths containing table t_i , we calculate the table score as: $\text{Score_Table}(t_i) = \max_{t \in \text{Path}_{t_i}} \text{Score_Path}(t)$. Finally, we select $T_N^Q = \{t_1, \dots, t_N\}$ with the highest $\text{Score_Table}(t_i)$ as our retrieval results. The detailed algorithm is provided in Appendix E. Appendix F elaborates on enhancing the table scoring algorithm to address ambiguous entities or synonyms in user questions or tables.

Dataset	#Table				
	1	2	3	4	All
SpiderUnion	395	214	43	6	658
BirdUnion	364	943	207	20	1534

Table 1: The distribution of questions based on the number of relevant tables (**#Table**) on the SpiderUnion and BirdUnion. **All** refers to the total number of questions in the dataset.

4 Experiments

4.1 Experiment Setup

Dataset To evaluate the effectiveness of MURRE, we validate it on two open-domain text-to-SQL datasets: SpiderUnion (Kothiyari et al., 2023) and BirdUnion, which combine the tables of Spider (Yu et al., 2018) and Bird (Li et al., 2023b), following previous works (Kothiyari et al., 2023; Chen et al., 2024b). Spider is a widely used cross-domain dataset for the text-to-SQL task, while Bird is more reflective of real-world scenarios, featuring more complicated queries. Table 1 presents the distribution of questions based on the number of tables required. Further details on Spider and Bird datasets are provided in Appendix G.

Metric We employ recall and complete recall as evaluation metrics for retrieval, and Execution Accuracy (EX) (Yu et al., 2018) for text-to-SQL. Recall ($r@$) measures the proportion of relevant tables retrieved from all relevant tables, following previous work (Kothiyari et al., 2023). Unlike other open-domain tasks (e.g., open-domain QA), retrieving all relevant tables is critical for generating correct SQL in open-domain text-to-SQL. Therefore, we introduce complete recall ($k =$), which is the proportion of examples where all relevant tables are retrieved. For text-to-SQL, we follow (Gao et al., 2023a) in using Execution Accuracy (EX) to assess the correctness of execution results by comparing the predicted SQL to the gold standard.

Model We utilize SGPT (Muennighoff, 2022) to embed tables and user questions without additional fine-tuning, following the previous work (Kothiyari et al., 2023). For the Removal phase and SQL generation, we use gpt-3.5-turbo³ to predict. The detailed descriptions of SGPT and gpt-3.5-turbo are provided in Appendix H.

³Document for gpt-3.5-turbo

Dataset	Model	Method	$k = 3$	$k = 5$	$k = 10$	$k = 20$	$r@3$	$r@5$	$r@10$	$r@20$
SpiderUnion	SGPT-125M	Single-hop	54.3	66.0	75.4	82.2	63.0	73.1	80.7	86.3
		CRUSH [†]	60.2	71.3	80.7	86.8	68.9	76.3	83.4	88.9
		MURRE	65.0	74.2	81.0	85.3	70.2	77.5	82.3	86.9
	SGPT-5.8B	Single-hop	76.3	86.8	94.1	97.6	84.0	91.5	96.2	98.7
		CRUSH [†]	68.2	80.1	88.4	92.2	75.5	85.1	91.2	94.5
		MURRE	86.0	93.5	96.7	97.3	89.3	94.3	96.8	97.5
BirdUnion	SGPT-125M	Single-hop	39.0	50.3	62.1	70.9	54.0	63.2	73.3	80.9
		CRUSH [†]	42.1	56.1	70.2	77.7	60.2	70.0	79.5	86.1
		MURRE	51.4	62.7	72.9	78.3	64.8	72.7	79.6	84.2
	SGPT-5.8B	Single-hop	55.3	67.3	79.4	86.4	72.9	80.8	88.6	92.8
		CRUSH [†]	52.2	63.5	78.4	88.1	70.0	77.9	87.5	93.0
		MURRE	69.1	80.1	88.7	92.7	81.0	87.6	92.6	95.4

Table 2: The complete recall (k) and recall (r) of MURRE, compared with Single-hop and CRUSH on SpiderUnion and BirdUnion, using SGPT-125M and SGPT-5.8B as the embedding models. [†] denotes our run since the performance difference led by the API change. The best results of each dataset and model are annotated in **bold**.

Comparing System In our experiments, we compare MURRE with the following methods: (i) Single-hop, which retrieves tables based on the user question in a single hop; (ii) CRUSH (Kothiyari et al., 2023), which retrieves tables in a single hop by converting the user question into a table format through hallucination.

Implement Details We set the beam size to 5, as it provides the best performance with the smallest size (see §4.5.1). The maximum hop (abbreviated as max hop) is set to 3 because over 98% of questions in SpiderUnion and BirdUnion require ≤ 3 tables (see Table 1). We use a 9-shot prompt for SpiderUnion and an 8-shot prompt for BirdUnion to remove the retrieved information, given the larger table scales in BirdUnion compared to SpiderUnion.

4.2 Main Result

The main results of our experiments are presented in Table 2. Compared to CRUSH, MURRE demonstrates significant improvements across datasets and models of different scales, with an average increase of 5.7% in recall and complete recall over the previous SOTA, which validates the effectiveness of our method. We can also see that:

The improvement of MURRE on BirdUnion is more significant than on SpiderUnion. Since the questions in BirdUnion typically require more tables (see Table 1), requiring multi-hop retrieval of MURRE more to obtain multiple relevant tables, thereby enhancing retrieval performance.

As the number of top-ranked tables grows, the retrieval performance of MURRE slows down.

Improving metrics with a high number of top-ranked tables necessitates retrieving relevant tables that are highly dissimilar to the user question, making the metrics challenging to enhance. Especially, for some metrics (e.g., $k = 20$, $r@20$), the performance of MURRE declines slightly because removing retrieved information shifts the focus toward retrieving increasingly dissimilar tables.

Text-to-SQL Experiments We conduct text-to-SQL experiments on SpiderUnion and BirdUnion using the user question and retrieved tables as the input, as shown in Table 3. The performance of MURRE exceeds that of both Single-hop and CRUSH, validating the effectiveness of our method. As the number of input tables increases, EX improvements slow beyond 10 tables because too many irrelevant tables hinder the model from focusing on the relevant ones. This also underscores the necessity of MURRE in enhancing the retrieval performance with a limited set of top-ranked tables in the open-domain text-to-SQL task.

4.3 Ablation Studies

To demonstrate the effectiveness of our method, we conduct ablation experiments on SpiderUnion, with results presented in Table 4. We select SpiderUnion corresponding to Spider for subsequent experiments because Spider is the mainstream text-to-SQL dataset. Since SGPT-125M and SGPT-5.8B exhibit similar trends across datasets and methods (see Table 2 and 3), we select SGPT-125M as the embedding model employed for subsequent experiments, considering both the embedding speed and the retrieval recall (Muennighoff et al., 2023).

Model	Method	SpiderUnion				BirdUnion			
		$r@3$	$r@5$	$r@10$	$r@20$	$r@3$	$r@5$	$r@10$	$r@20$
SGPT-125M	Single-hop	43.9	50.0	53.2	54.1	11.2	13.4	17.1	18.5
	CRUSH	47.3	50.9	55.9	59.6	14.5	17.1	19.0	20.5
	MURRE	50.3	54.1	54.7	57.4	16.0	16.8	19.4	20.0
SGPT-5.8B	Single-hop	54.9	60.6	61.6	63.7	16.9	17.3	18.3	20.4
	CRUSH	49.8	56.4	60.3	60.8	16.8	18.0	19.7	21.1
	MURRE	62.5	64.4	64.4	66.7	20.9	21.8	22.0	22.4

Table 3: EX for predicted SQL based on the user question and varying numbers of retrieved tables.

Method	$k = 3$	$k = 5$	$k = 10$	$r@3$	$r@5$	$r@10$
MURRE	65.0	74.2	81.0	70.2	77.5	82.3
<i>w/o removal</i>	46.2 (-18.8)	56.7 (-17.5)	67.2 (-13.8)	50.6 (-19.6)	60.7 (-16.8)	70.0 (-11.6)
<i>w/o tabulation</i>	54.6 (-10.4)	64.9 (-9.3)	75.5 (-5.5)	63.4 (-6.8)	72.5 (-5.0)	80.9 (-1.4)
<i>w/o early stop</i>	52.6 (-12.4)	64.9 (-9.3)	71.0 (-10.0)	57.1 (-13.1)	67.0 (-10.5)	72.2 (-10.1)

Table 4: The ablation results on evaluating MURRE, compared with splicing the question with retrieved tables (denoted as *w/o removal*), querying the rest information with natural language after removing retrieved information (denoted as *w/o tabulation*), and without employing the mechanism of early stop (denoted as *w/o early stop*) on SpiderUnion with SGPT-125M.

The Effectiveness of Removal To demonstrate the effectiveness of Removal which removes retrieved information from the user question, we compare its performance against the popular multi-hop retrieval method in open-domain QA, which directly splices the user question with retrieved tables at each hop. Compared with MURRE, the performance of the splicing method drops significantly, underscoring the effectiveness of Removal which alleviates retrieving similar but irrelevant tables.

The Effectiveness of Tabulation To prove the effectiveness of transforming questions into a tabular format (abbreviated as tabulation), we query the removed questions with natural language. The results indicate that, compared to querying with natural language, tabulation significantly enhances performance, validating its effectiveness in MURRE. The prompt without tabulation is presented in Appendix B.

The Effectiveness of Early Stop To verify the effectiveness of early stop in MURRE, we compare the results without using this mechanism, where the model does not generate the special early stop mark. The performance without the early stop shows a significant decline, demonstrating that its inclusion in MURRE is crucial for maintaining performance.

4.4 Analysis

In the analysis experiments, we use $k = 5$ as the evaluation metric, with detailed explanations in Appendix I. Additionally, we discuss the efficiency of




Figure 3: The proportion of performance improvements achieved by MURRE in addressing various limitations, compared to Single-hop. The numbers in parentheses in the legend represent the number of improved examples.

MURRE in Appendix J, the reasons we use the original user question at the first hop in Appendix K, and the impact of SQL hardness in Appendix L.

4.5 Why MURRE improve compared with Single-hop method?

To explore why MURRE enhances retrieval performance, we analyze its improvement compared with Single-hop in Figure 3. We observe that: (i) Similar Candidate means that many irrelevant candidate tables are similar to the question. MURRE alleviates the limitation by removing the retrieved information, allowing the model to focus on retrieving tables not yet retrieved. (ii) Semantic Gap refers to the semantic distinction between the natural language question and the relevant tables. MURRE tabulates the question concerning the retrieved information, aligning the question more closely with tables in terms of domain and vocabulary, thus narrowing the semantic gap (Martino et al., 2023; Li

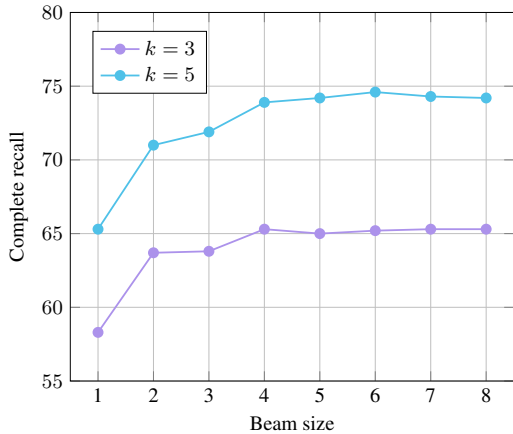


Figure 4: The complete recall with different beam sizes on SpiderUnion with SGPT-125M.

Max Hop	#Table				
	1	2	3	≥ 4	All
1	73.7	59.8	25.6	50.0	66.0
2	73.2	77.6	58.1	50.0	73.4
3	74.2	78.0	58.1	50.0	74.2
4	74.2	78.0	58.1	50.0	74.2

Table 5: Complete recall $k = 5$ of MURRE with varying max hops. SpiderUnion is divided based on the number of relevant tables per question (denoted as #Table). All represents the entire SpiderUnion dataset. The best results for each #table division are annotated in **bold**.

et al., 2024). Detailed examples illustrating the advantages of MURRE are provided in Appendix M, with the statistical criteria in Appendix N.

4.5.1 How Does Beam Size Affect the Performance?

To observe the impact of different beam sizes on the retrieval performance, we compare the performance of MURRE with different beam sizes, as shown in Figure 4. (i) When MURRE does not employ beam search, i.e., with a beam size of 1, performance degrades rapidly, suggesting that beam search mitigates the effects of error cascade. (ii) As the beam size increases, complete recall improves significantly until a beam size of 5, after which performance either plateaus or declines. This indicates that while employing beam search enhances performance, a beam size ≤ 5 introduces too many irrelevant tables, increasing computational cost without further performance improvement.

4.5.2 How Does the Number of Hops Affect the Performance?

To verify the effectiveness of multi-hop retrieval in MURRE, we compare the performance with varying the number of max hops, as shown in Table 5.

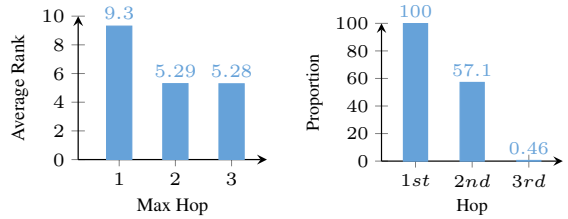


Figure 5: The left part is the average rank of relevant tables with different numbers of max hops on the SpiderUnion using MURRE. The right part is the proportion of questions that are not early stopped with different hops on the SpiderUnion using MURRE.

The results indicate: (i) MURRE performs optimally when the number of max hop is greater than or equal to the number of required tables. (ii) For questions requiring 1 or 2 tables, the best performance is achieved at a max hop of 3. Since MURRE can enhance performance for the questions that retrieve irrelevant tables by guiding to obtain unretrieved relevant tables in subsequent hops. (iii) For questions requiring 1 table, performance slightly decreases because removing retrieved information from such questions could easily introduce errors. (iv) The performance of requiring ≥ 4 tables remains unchanged because improving complete recall $k = 5$ requires all relevant tables included in the top 5 retrieved, which is challenging.

4.5.3 Can MURRE Reduce the Average Rank of Relevant Tables?

To verify MURRE can remove the retrieved information from the question and retrieve relevant tables that are not similar to previously retrieved tables, we calculate the average rank of relevant tables at different maximum hops, as shown in the left part of Figure 5. The figure presents that: (i) MURRE significantly enhances the average rank of relevant tables, with the most notable improvement occurring at a maximum hop of 2. This is because most questions in SpiderUnion require 1 or 2 tables (see Table 1), requiring two hops to obtain unretrieved relevant tables. (ii) Conversely, the improvement at a maximum hop of 3 is weak, as only a small number of questions require ≥ 3 tables but also because, and most questions complete retrieval before the third hop.

As illustrated in the right part of Figure 5, the proportion of retrieving at different hops in MURRE closely matches the distribution of the number of relevant tables in Table 1. We set the maximum hop to 3 also because most user questions stop retrieving in the third hop.

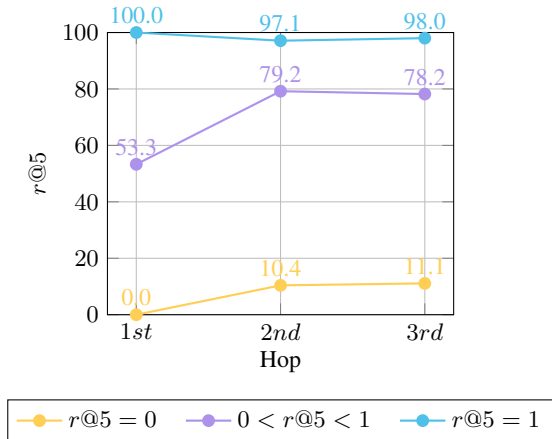


Figure 6: The $r@5$ during hops, categorizing SpiderUnion according to the $r@5$ in the first hop which falls into different intervals.

4.5.4 How Does the Previous Errors Affect Subsequent Performance?

To analyze the error cascades in MURRE, we divide questions based on first-hop results: $r@5 = 0$, $0 < r@5 < 1$, and $r@5 = 1$, and compare performance across hops, as shown in Figure 6. The error cascades in MURRE are minimal and significantly compensated by the performance improvements.

We observe that: (i) For questions whose $r@5 = 0$ in the first hop, all the top 5 tables retrieved are irrelevant. MURRE improves performance by removing retrieved information from the questions, thereby eliminating interference from irrelevant tables. (ii) For questions with $0 < r@5 < 1$ in the first hop, some of the top 5 tables are irrelevant. MURRE significantly enhances performance by removing the retrieved information from the user question and focusing on unretrieved information. Performance in the third hop slightly declines mainly because most questions in SpiderUnion involve ≤ 2 relevant tables, and additional hops may introduce errors. (iii) For questions whose $r@5 = 1$ in the first hop, performance in subsequent hops is slightly reduced as all relevant tables are retrieved, causing subsequent hops to introduce minor errors.

5 Related Work

5.1 Text-to-SQL

The text-to-SQL task aims to convert user questions into SQL queries, facilitating efficient database access (Qin et al., 2022). LLM-based methods have become mainstream in text-to-SQL due to their superior performance with minimal annotated data (Li et al., 2023a; Gao et al., 2023a). For ex-

ample, Li and Xie (2024) propose using LLMs to predict execution results of test cases, determining the correctness of candidates SQL queries. However, most previous methods do not focus on open-domain text-to-SQL, and exist a gap with real-world applications. Therefore, CRUSH (Kothyari et al., 2023) uses LLMs to guess potentially relevant tables for retrieval. DBCopilot (Wang et al., 2024) trains a schema router to identify relevant tables. Chen et al. (2024b) propose a re-ranking relevance method by fine-tuning DTR models.

However, existing methods are constrained by: (i) single-hop retrieval; (ii) requiring fine-tuning, which is resource-intensive and domain-specific. To solve these problems, we propose a multi-hop table retrieval method for open-domain text-to-SQL.

5.2 Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) denotes retrieving information in extra knowledge bases to support accurate and reliable LLM generation (Chen et al., 2024a). Existing RAG works employ either single-hop retrieval (Ma et al., 2023; Zheng et al., 2024) or multi-hop retrieval, where questions in each hop can be generated in two ways (Gao et al., 2023b): (i) splicing the question with the current and previous retrieval results (Shao et al., 2023; Jeong et al., 2024), or (ii) using LLMs to integrate the retrieved results into the question (Press et al., 2023; Trivedi et al., 2023; Li et al., 2023c).

To enhance retrieval in open-domain text-to-SQL, we adopt multi-hop retrieval. However, previous methods are unsuitable for this task, because adding retrieved tables to the question could lead to the retrieval of similar yet irrelevant tables, unlike other tasks, where questions require supplementary. Therefore, we remove retrieved information from the question for next-hop retrieval.

6 Conclusion

In the paper, we propose MURRE to address the challenge that multi-hop retrieval in other open-domain tasks cannot be directly applied to open-domain text-to-SQL. Compared with previous methods, MURRE removes the retrieved information from the question to ensure the retrieved tables are relevant to the question. Experimental results demonstrate the effectiveness of MURRE on two open-domain text-to-SQL datasets. Our method achieves new SOTA results compared to previous methods, with an average of 5.7% improvement.

Limitations

We discuss the limitations of our work from the following two aspects. (i) Considering the applicability, the multi-turn text-to-SQL task is common in real scenarios (Yu et al., 2019a,b), while we do not discuss the solutions of open-domain multi-turn text-to-SQL. We leave improving our method to apply to multi-turn text-to-SQL for future work. (ii) From the performance perspective, our method does not consider the performance improvement brought by the text-to-SQL feedback (Trivedi et al., 2023; Yu et al., 2023). We leave the retrieval recall improvement leveraging the results of text-to-SQL for future work. handles ambiguous entities or synonyms within the natural language questions or database schemas. Although our method achieves significant improvements, future work can improve our method from the aspects of applicability and recall further.

Ethics Statement

Every dataset and model used in the paper is accessible to the public, and our application of them adheres to their respective licenses and conditions.

Acknowledgment

We gratefully acknowledge the support of the National Natural Science Foundation of China (NSFC) via grant 62236004, 62206078, 62441603 and 62476073.

References

- Jiawei Chen, Hongyu Lin, Xianpei Han, and Le Sun. 2024a. [Benchmarking large language models in retrieval-augmented generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17754–17762.
- Peter Baile Chen, Yi Zhang, and Dan Roth. 2024b. [Is table retrieval a solved problem? exploring join-aware multi-table retrieval](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2687–2699, Bangkok, Thailand. Association for Computational Linguistics.
- Longxu Dou, Yan Gao, Mingyang Pan, Dingzirui Wang, Wanxiang Che, Jian-Guang Lou, and Dechen Zhan. 2023. Unisar: A unified structure-aware autoregressive language model for text-to-sql semantic parsing. *International Journal of Machine Learning and Cybernetics*, 14(12):4361–4376.
- Yair Feldman and Ran El-Yaniv. 2019. [Multi-hop paragraph retrieval for open-domain question answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2296–2309, Florence, Italy. Association for Computational Linguistics.
- Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. 2023a. [Text-to-sql empowered by large language models: A benchmark evaluation](#). *ArXiv*, abs/2308.15363.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2023b. [Retrieval-augmented generation for large language models: A survey](#). *ArXiv*, abs/2312.10997.
- Kazuma Hashimoto, Iftexhar Naim, and Karthik Ramman. 2024. [How does beam search improve span-level confidence estimation in generative sequence labeling?](#) In *Proceedings of the 1st Workshop on Uncertainty-Aware NLP (UncertainNLP 2024)*, pages 62–69, St Julians, Malta. Association for Computational Linguistics.
- Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong Park. 2024. [Adaptive-rag: Learning to adapt retrieval-augmented large language models through question complexity](#). In *NAACL*.
- Chris Kamphuis, Arjen P. de Vries, Leonid Boytsov, and Jimmy J. Lin. 2020. [Which bm25 do you mean? a large-scale reproducibility study of scoring variants](#). *Advances in Information Retrieval*, 12036:28 – 34.
- Kezhi Kong, Jiani Zhang, Zhengyuan Shen, Balasubramaniam Srinivasan, Chuan Lei, Christos Faloutsos, Huzefa Rangwala, and George Karypis. 2024. [Opentab: Advancing large language models as open-domain table reasoners](#). In *The Twelfth International Conference on Learning Representations*.
- Mayank Kothiyari, Dhruva Dhingra, Sunita Sarawagi, and Soumen Chakrabarti. 2023. [CRUSH4SQL: Collective retrieval using schema hallucination for Text2SQL](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Hyunji Lee, Sohee Yang, Hanseok Oh, and Minjoon Seo. 2022. [Generative multi-hop retrieval](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. 2023a. [Resdsq: decoupling schema linking and skeleton parsing for text-to-sql](#). In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*. AAAI Press.
- Jiarui Li, Ye Yuan, and Zehua Zhang. 2024. [Enhancing llm factual accuracy with rag to counter hallucinations: A case study on domain-specific queries in private knowledge-bases](#). *Preprint*, arXiv:2403.10446.

- Jinyang Li, Binyuan Hui, GE QU, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin Chang, Fei Huang, Reynold Cheng, and Yongbin Li. 2023b. [Can LLM already serve as a database interface? a BIG bench for large-scale database grounded text-to-SQLs](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Xiaonan Li, Changtai Zhu, Linyang Li, Zhangyue Yin, Tianxiang Sun, and Xipeng Qiu. 2023c. [Llatrieval: Llm-verified retrieval for verifiable generation](#). *ArXiv*, abs/2311.07838.
- Zhenwen Li and Tao Xie. 2024. [Using llm to select the right sql query from candidates](#). *Preprint*, arXiv:2401.02115.
- Qian Liu, Dejian Yang, Jiahui Zhang, Jiaqi Guo, Bin Zhou, and Jian-Guang Lou. 2021. [Awakening latent grounding from pretrained language models for semantic parsing](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1174–1189, Online. Association for Computational Linguistics.
- Xinbei Ma, Yeyun Gong, Pengcheng He, Hai Zhao, and Nan Duan. 2023. [Query rewriting in retrieval-augmented large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5303–5315, Singapore. Association for Computational Linguistics.
- Ariana Martino, Michael Iannelli, and Coleen Truong. 2023. [Knowledge injection to counter large language model \(llm\) hallucination](#). In *The Semantic Web: ESWC 2023 Satellite Events*, pages 182–185, Cham. Springer Nature Switzerland.
- Niklas Muennighoff. 2022. [Sgpt: Gpt sentence embeddings for semantic search](#). *ArXiv*, abs/2202.08904.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive text embedding benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. [Measuring and narrowing the compositionality gap in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Association for Computational Linguistics.
- Bowen Qin, Binyuan Hui, Lihan Wang, Min Yang, Jinyang Li, Binhua Li, Ruiying Geng, Rongyu Cao, Jian Sun, Luo Si, Fei Huang, and Yongbin Li. 2022. [A survey on text-to-sql parsing: Concepts, methods, and future directions](#). *ArXiv*, abs/2208.13629.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. [Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*. Association for Computational Linguistics.
- Liang Shi, Zhengju Tang, Nan Zhang, Xiaotong Zhang, and Zhi Yang. 2024. [A survey on employing large language models for text-to-sql tasks](#). *Preprint*, arXiv:2407.15186.
- Harsh Trivedi, Niranjana Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Shuai Wang, Shengyao Zhuang, and Guido Zuccon. 2021. [Bert-based dense retrievers require interpolation with bm25 for effective passage retrieval](#). In *Proceedings of the 2021 ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR '21*. Association for Computing Machinery.
- Tianshu Wang, Hongyu Lin, Xianpei Han, Le Sun, Xiaoyang Chen, Hao Wang, and Zhenyu Zeng. 2024. [Dbcopilot: Scaling natural language querying to massive databases](#). *Preprint*, arXiv:2312.03463.
- Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, Xu Zhao, Min-Yen Kan, Junxian He, and Qizhe Xie. 2023. [Self-evaluation guided beam search for reasoning](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Wenhan Xiong, Xiang Li, Srini Iyer, Jingfei Du, Patrick Lewis, William Yang Wang, Yashar Mehdad, Scott Yih, Sebastian Riedel, Douwe Kiela, and Barlas Oguz. 2021. [Answering complex open-domain questions with multi-hop dense retrieval](#). In *International Conference on Learning Representations*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik R Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Tao Yu, Rui Zhang, Heyang Er, Suyi Li, Eric Xue, Bo Pang, Xi Victoria Lin, Yi Chern Tan, Tianze Shi, Zihan Li, Youxuan Jiang, Michihiro Yasunaga, Sungrok Shim, Tao Chen, Alexander Fabbri, Zifan Li, Luyao Chen, Yuwen Zhang, Shreya Dixit, Vincent Zhang, Caiming Xiong, Richard Socher, Walter Lasecki, and Dragomir Radev. 2019a. [CoSQL: A](#)

- conversational text-to-SQL challenge towards cross-domain natural language interfaces to databases. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1962–1979, Hong Kong, China. Association for Computational Linguistics.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 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*. Association for Computational Linguistics.
- Tao Yu, Rui Zhang, Michihiro Yasunaga, Yi Chern Tan, Xi Victoria Lin, Suyi Li, Heyang Er, Irene Li, Bo Pang, Tao Chen, Emily Ji, Shreya Dixit, David Proctor, Sungrok Shim, Jonathan Kraft, Vincent Zhang, Caiming Xiong, Richard Socher, and Dragomir Radev. 2019b. [SPaRC: Cross-domain semantic parsing in context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4511–4523, Florence, Italy. Association for Computational Linguistics.
- Zichun Yu, Chenyan Xiong, Shi Yu, and Zhiyuan Liu. 2023. [Augmentation-adapted retriever improves generalization of language models as generic plug-in](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Jiahao Zhang, Haiyang Zhang, Dongmei Zhang, Yong Liu, and Shen Huang. 2024. [End-to-end beam retrieval for multi-hop question answering](#). In *2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics*.
- Wayne Xin Zhao, Jing Liu, Ruiyang Ren, and Ji-Rong Wen. 2023a. [Dense text retrieval based on pretrained language models: A survey](#). *ACM Trans. Inf. Syst.* Just Accepted.
- 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. 2023b. [A survey of large language models](#). *Preprint*, arXiv:2303.18223.
- Huaxiu Steven Zheng, Swaroop Mishra, Xinyun Chen, Heng-Tze Cheng, Ed H. Chi, Quoc V Le, and Denny Zhou. 2024. [Take a step back: Evoking reasoning via abstraction in large language models](#). In *The Twelfth International Conference on Learning Representations*.

A How to Address the Scenarios Where No Relevant Tables Exist

In this section, we discuss how our method can be improved to solve the situation that there are no tables that are relevant to the user question in the given databases (Yu et al., 2019a). According to the existing open-domain database tables, we can manually annotate or synthesize relevant and irrelevant questions to train a discriminator. For each user question, before retrieving the relevant tables, we can use the discriminator to determine whether the question is irrelevant to the existing tables. And if the question is irrelevant, we can directly output the feedback, and no longer retrieve tables and generate SQL (Jeong et al., 2024).

B Prompts for Removal

In the section, we show the prompts we use to remove retrieved information from the question on SpiderUnion (see Table 6) and BirdUnion (see Table 7). We also show the prompt without tabulation which is used in §4.3, as shown in Table 8. Each table is represented in the form of “database name.table name(column name, column name, ...)” following Kothiyari et al. (2023). We only show the first two examples here limited by pages. The code and the whole prompt will be public in the future.

C Normalization Method

In this section, we show how to normalize the cosine similarity into a probability distributed between 0 and 1 present in Equation 3.1. We define the cosine similarity between the question q vector and the table t_i vector as s , which is distributed between -1 and 1 . And we use Equation C.1 to normalize the cosine similarity s .

$$Norm(s) = \frac{s + 1}{2} \quad (C.1)$$

Moreover, $Norm(s)$ is proportional to s , that is, the greater the cosine similarity s , the greater $\hat{P}(t_i|q)$, that is, the greater the probability that the table t_i is retrieved by q .

D Score of the Retrieval Path

In this section, we prove the calculation process of the retrieval path probability present in §3.5. First of all, we define the retrieval path $\text{Path}^{h,b}$ as Equation D.1, where $q^{h,b}$ represents the user

question of hop h and beam b , and $t_{p_h}^{q^{h,b}}$ represents the table retrieved by $q^{h,b}$ ranked at p_h .

$$\text{Path}^{h,b} = ((q^{1,b}, t_{p_1}^{q^{1,b}}), \dots, (q^{h,b}, t_{p_h}^{q^{h,b}})) \quad (D.1)$$

According to the discussion in §3.3, the score of each retrieval path $\text{Path}^{h,b}$ can be regarded as the probability that the last table in the path $t_{p_h}^{q^{h,b}}$ is retrieved in the case of the user question at last hop $q^{h,b}$. Formally, it can be summarized as:

$$\begin{aligned} & \text{Score_Path}((q^{1,b}, t_{p_1}^{q^{1,b}}), \dots, (q^{h,b}, t_{p_h}^{q^{h,b}})) \\ &= \hat{P}((q^{1,b}, t_{p_1}^{q^{1,b}}), \dots, (q^{h,b}, t_{p_h}^{q^{h,b}})) \\ &= \hat{P}(t_{p_h}^{q^{h,b}} | q^{h,b}) \cdot \hat{P}((q^{1,b}, t_{p_1}^{q^{1,b}}), \dots, (q^{h-1,b}, t_{p_{h-1}}^{q^{h-1,b}})) \\ &= \dots \\ &= \prod_{j=1}^h \hat{P}(t_{p_j}^{q^{j,b}} | q^{j,b}) \end{aligned} \quad (D.2)$$

Therefore, we multiply all \hat{P} on the retrieval path as the score of the path.

E Table Scoring Algorithm in MURRE

In this section, we detail the table scoring algorithm (Algorithm 1), which is discussed in §3.5.

F How to Handle Ambiguous Entities or Synonyms

In this section, we present how our method can be improved to handle ambiguous entities or synonyms within user questions or database tables. When scoring the retrieved tables, in the face of multiple similar tables retrieved due to ambiguous entities or synonyms, we can additionally consider the correlation between tables and select tables with higher relevance scores of other retrieved tables (Chen et al., 2024b).

G Dataset Details

In this section, we introduce in detail the source dataset of SpiderUnion and BirdUnion datasets which we use. Spider (Yu et al., 2018) is a multi-domain mainstream text-to-SQL dataset that contains 658 questions, with an average of 1.48 tables per question in the dev-set. Bird (Li et al., 2023b), as a text-to-SQL dataset, is closer to the actual scenario featuring its larger scale and more difficult questions. Bird contains 1534 questions, with an average of 1.92 tables per question in the dev-set.

Given the following SQL tables, your job is to complete the possible left SQL tables given a user’s request. Return None if no left SQL tables according to the user’s request.

Question: Which models are lighter than 3500 but not built by the ‘Ford Motor Company’?
Database: car_1.model list(model id, maker, model)
car_1.cars data(id, mpg, cylinders, edispl, horsepower, weight, accelerate, year)
car_1.car names(make id, model, make)
Completing Tables: car_1.car makers(maker)

Question: Which employee received the biggest bonus? Give me the employee name.
Database: employee_hire_evaluation.evaluation(employee id, year awarded, bonus)
employee_hire_evaluation.employee(employee id, name, age, city)
Completing Tables: None

...

Table 6: The prompt we use for the SpiderUnion with gpt-3.5-turbo.

Given the following SQL tables, your job is to complete the possible left SQL tables given a user’s request. Return None if no left SQL tables according to the user’s request.

Question: What was the growth rate of the total amount of loans across all accounts for a male client between 1996 and 1997?
Database: financial.client(client_id, gender, birth_date, location of branch)
financial.loan(loan_id, account_id, date, amount, duration, monthly payments, status)
Completing Tables: financial.account(account id, location of branch, frequency, date)
financial.disp(disposition id, client_id, account_id, type)

Question: How many members did attend the event ‘Community Theater’ in 2019?
Database: student_club.Attendance(link to event, link to member)
Completing Tables: student_club.Event(event name, event date)

...

Table 7: The prompt we use for the BirdUnion with gpt-3.5-turbo.

H Model Details

In the section, we introduce the models SGPT and gpt-3.5-turbo used in our experiments. SGPT (Muennighoff, 2022) is the popular retrieval Single-hop, employing a decoder-only architecture and showing excellent performance on tasks such as sentence matching. gpt-3.5-turbo (Zhao et al., 2023b) has undergone instruction fine-tuning and human alignment and has superior in-context learning and inference capability.

I The Evaluation Metric in Analysis Experiments

In this section, we explain the reasons for using complete recall $k = 5$ as the evaluation metric in the analysis experiments. The increasing trend of the performance in the text-to-SQL becomes slow or even drops when inputting retrieved tables more than 5 as shown in Table 3, and considering that SpiderUnion and BirdUnion require up to 4 tables for each question, so in the following analysis, we are mainly concerned with the performance of the top 5 retrieval results. Furthermore, complete recall $k = 5$ is a more strict indicator than $recall@5$, so

we mainly utilize complete recall $k = 5$ as the evaluation metric in analysis experiments.

J Discussion on Efficiency

In this section, we discuss the comparison of efficiency between MURRE and CRUSH. Because of each user question, CRUSH needs to use LLM to predict the relevant tables once, and then retrieve all the tables according to the LLM prediction once. So the time complexity of CRUSH is shown in Equation J.1, where n is the number of user questions.

$$\begin{aligned} T(CRUSH) &= O(2 \cdot n) \\ &= O(n) \end{aligned} \tag{J.1}$$

Suppose that the number of hop is H and the beam size is B in MURRE. For each user question, MURRE needs to retrieve all the tables first, and input LLM for removal according to the retrieved top B tables. In the subsequent hops, each hop needs to retrieve B times and remove information B times with LLM. Therefore, the time complexity

Remove information appearing in the database from the question .
Return None if the database is totally correspond to the question .

Question: Which models are lighter than 3500 but not built by the 'Ford Motor Company'?
Database: car_1.model list(model id, maker, model)
car_1.cars data(id, mpg, cylinders, edispl, horsepower, weight, accelerate, year)
car_1.car names(make id, model, make)
Rewritten Question: What is the car makers of the 'Ford Motor Company'?

Question: Which employee received the biggest bonus? Give me the employee name.
Database: employee_hire_evaluation.evaluation(employee id, year awarded, bonus)
employee_hire_evaluation.employee(employee id, name, age, city)
Rewritten Question: None

...

Table 8: The prompt we use for the SpiderUnion with gpt-3.5-turbo **without tabulation**.

Algorithm 1 The **table scoring** algorithm in MURRE

Input: The similarity corresponding to each table t in each hop h : $all_paths = [[(table_{11}, score_{11}), \dots, (table_{1H}, score_{1H})], \dots, [(table_{P1}, score_{P1}), \dots, (table_{PH}, score_{PH})]]$, the number of max hops H , the number of all paths P .

Output: The scores of each table t

- 1: Initialization : $table_score \leftarrow \{\}$
- 2: **for** $each_path$ in all_paths **do**
- 3: $score \leftarrow 1$ ▷ Initialize the score
- 4: **for** $example$ in $each_path$ **do**
- 5: $score = score \times example[1]$ ▷ Calculate the score of the path
- 6: **end for**
- 7: **for** $example$ in $each_path$ **do**
- 8: $table_score[example[0]] \leftarrow max(score, table_score[example[0]])$ ▷ Update the table score with the max path score
- 9: **end for**
- 10: **end for**
- 11: **return** $table_score$

of our method is present in Equation J.2.

$$\begin{aligned}
T(\text{MURRE}) &= O((1 + B + (H - 1) \cdot B \cdot 2) \cdot n) \\
&= O(B \cdot H \cdot 2 \cdot n) \\
&= O((B \cdot H) \cdot n)
\end{aligned}
\tag{J.2}$$

It can be found that although our method has significantly improved the performance compared with CRUSH, our method is less efficient. However, existing work shows that using reasoning efficiency for improving the reasoning performance has a wide range of practical application value (Yao et al., 2023; Xie et al., 2023; Press et al., 2023; Hashimoto et al., 2024). Therefore, in practical applications, how to choose B and H in MURRE to achieve a balance between retrieval efficiency and effect should be carefully considered.

K Why Retrieve with the Original Question at the First Hop

To demonstrate the effectiveness of MURRE which employs the original user question at the first hop rather than the tabulation question, we conduct experiments, with the results presented in Table 9. The results show that starting with tabulation question in table form reduces the retrieval performance, thereby validating the effectiveness of MURRE. The decline in performance is attributed to the absence of the retrieved table to provide domain-constraining information, which increases errors when the question is tabulated at the first hop.

L Impact of SQL Hardness

In this section, we show the performance of MURRE on SQL of different hardness levels. We categorize the SQL and its corresponding question

Dataset	Method	$k = 3$	$k = 5$	$k = 10$	$k = 20$	$r@3$	$r@5$	$r@10$	$r@20$
BIRDUunion	MURRE	69.1	80.1	88.7	92.7	81.0	87.6	92.6	95.4
	<i>w tabulation at the first hop</i>	52.2	63.6	78.4	88.1	70.0	78.0	87.5	93.0
SpiderUnion	MURRE	86.0	93.5	96.7	97.3	89.3	94.3	96.8	97.5
	<i>w tabulation at the first hop</i>	69.3	80.2	88.3	92.2	76.1	85.1	91.1	94.5

Table 9: The comparison between MURRE and the method of tabulating the question at the first hop, with using SGPT-5.8B as the embedding and using gpt-3.5-turbo to tabulate.

Method	Easy	Medium	Hard	Extra	All
Single-hop	70.5	71.1	55.8	51.3	66.0
MURRE	71.8	76.0	73.3	73.1	74.2

Table 10: Complete recall $k = 5$ of MURRE compared with the Single-hop in different SQL hardness levels on SpiderUnion. **Extra** denotes extra hard. **All** refers to the performance of the whole SpiderUnion dataset. The best results of different hardness are annotated in **bold**.

Question
What is the most populace city that speaks English?
Tables
city_record.city(city id, city, hanyu pinyin, regional population, ...)
world_1.city(id, name, country code, district, population)
e_government.addresses(address id, line 1 number building, ...)
...
Retrieved Tables (top 3)
CRUSH: ($r@3 = 50.0$)
farm.city(city id, official name, status, area km 2, population, ...)
world_1.city(id, name, country code, district, population)
geo.city(city name, population, country name, state name)
MURRE: ($r@3 = 100.0$)
world_1.city(id, name, country code, district, population)
world_1.countrylanguage(countrycode, language, is official, ...)
city_record.city(city id, city, hanyu pinyin, regional population, ...)

Figure 7: A case study comparing MURRE with CRUSH. Green indicates relevant tables, while red indicates irrelevant ones. Each table is represented as “database name.table name(column names)”. r denotes recall.

according to the SQL hardness criteria (Yu et al., 2018) and calculate the retrieval performance of different hardness levels, as shown in Table 10. MURRE improves performance more significantly for more difficult SQL questions. Because more difficult SQL often requires more tables to operate and query, the Single-hop is challenging to retrieve all relevant tables merely in a single hop, while our method can retrieve more relevant tables with multi-hop retrieval by removing the retrieved information from the question at each hop.

M Case Study

In this section, we compare our methods with other retrieval methods through specific examples to illustrate the advantages of our methods. Firstly, we demonstrate a case study with MURRE compared with CRUSH, as shown in Figure 7. We can see that CRUSH fails to retrieve the table “world_1.countrylanguage” within the top 3 results due to its single-hop retrieval limitation, as the retrieval of table “world_1.countrylanguage” relies on the table “world_1.city”. In contrast, MURRE employs multi-hop retrieval with a beam size of 3, increasing the probability of selecting relevant tables at each hop. Additionally, we eliminate the retrieved information in “world_1.city” from the question, which aids the model in retrieving “world_1.countrylanguage” that was previously missed.

We also present one example in detail comparing MURRE with the Single-hop, without beam search and without Removal respectively in Table 11, Table 12, and Table 13. We set the beam_size to 3, max hop to 3, and MURRE stops early at the second hop.

As shown in Table 11, the Single-hop retrieval fails to retrieve the table “world_1.countrylanguage” at top 3 limited by the single-hop retrieval since the retrieval of table “world_1.countrylanguage” relies on the table “world_1.city”. As displayed in Table 12, MURRE without beam search method is affected by error cascades, because the table city_record.city with the highest retrieval ranking in hop 1 is irrelevant to the question. Removing based on the irrelevant city_record.city table lead to retrieval errors in subsequent hops. As present in Table 13, MURRE without Removal adds the retrieved tables directly to the user question, so that the subsequent retrieved tables are similar to the currently retrieved tables. For example, the irrelevant table “city_record.hosting city” retrieved in hop 2 is similar to the table “city_record.city” retrieved in

Question

What is the most populace city that speaks English?

Single-hop ($r@3 = 50.0$)

Retrieved Tables (top 3)

`city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)`

`world_1.city(id, name, country code, district, population)`

`e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)`

MURRE ($r@3 = 100.0$)

Retrieved Tables, hop = 1 (top 3)

`city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)`

`world_1.city(id, name, country code, district, population)`

`e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)`

Removal Questions

`city_record.language(city id, language, percentage)`

`world_1.countrylanguage(countrycode, language, is official, percentage)`

`e_government.languages(language id, language name, language code, population)`

Retrieved Tables, hop = 2 (top 3)

`world_1.city(id, name, country code, district, population)`

`world_1.countrylanguage(countrycode, language, is official, percentage)`

`city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)`

Removal Questions

None

None

None

(Early Stop)

Table 11: Detailed case study comparing MURRE with Single-hop. The green means the relevant table, while the red means irrelevant. Each table is expressed in the form of “database name.table name(column names)”. r denotes recall.

hop 2, which are both about “city” information, but ignore the information of “language”. Our method focuses on retrieving tables about “language” by removing the information “world_1.city” in the retrieved tables, and successfully retrieves two relevant tables.

the relevant tables, the retrieval similarity is also low, representing the semantic gap to some extent (Reimers and Gurevych, 2019).

N Statistical Criteria of Limitations

To facilitate statistics on the number of results reflected in the two limitations of Similar Candidate and Semantic Gap, we set the following rules.

For the limitation of Similar Candidate, the statistical standard is that if the irrelevant table that is incorrectly retrieved has the same token as the question, it is considered a Similar Candidate error. If the same token appears in the schema as in the question, the cosine similarity between the table and the question after embedding is also high, making it challenging to identify the relevant tables (Kamphuis et al., 2020; Wang et al., 2021; Li et al., 2023c; Kong et al., 2024; Wang et al., 2024).

For the limitation of Semantic Gap, the statistical standard is that if the relevant table that has not been retrieved does not have the same token as the question, it is considered a Semantic Gap error. If the question does not overlap with any token in

Question

What is the most populace city that speaks English?

MURRE without beam search ($r@3 = 0.0$)

Retrieved Tables, hop = 1 (top 3)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

world_1.city(id, name, country code, district, population)

e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)

Removal Question

city_record.language(city id, language, percentage)

Retrieved Tables, hop = 2 (top 3)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

city_record.hosting city(year, match id, host city)

city_record.match(match id, date, venue, score, result, competition)

Removal Question

None

(Early Stop)

MURRE ($r@3 = 100.0$)

Retrieved Tables, hop = 1 (top 3)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

world_1.city(id, name, country code, district, population)

e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)

Removal Questions

city_record.language(city id, language, percentage)

world_1.countrylanguage(countrycode, language, is official, percentage)

e_government.languages(language id, language name, language code, population)

Retrieved Tables, hop = 2 (top 3)

world_1.city(id, name, country code, district, population))

world_1.countrylanguage(countrycode, language, is official, percentage)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

Removal Questions

None

None

None

(Early Stop)

Table 12: Detailed case study comparing MURRE with MURRE without beam search. The green means the relevant table, while the red means irrelevant. Each table is expressed in the form of “database.name.table name(column names)”. r denotes recall.

Question

What is the most populace city that speaks English?

MURRE without Removal ($r@3 = 50.0$)

Retrieved Tables, hop = 1 (top 3)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

world_1.city(id, name, country code, district, population)

e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)

Spliced Questions

What is the most populace city that speaks English?; city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

What is the most populace city that speaks English?; world_1.city(id, name, country code, district, population)

What is the most populace city that speaks English?; e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)

Retrieved Tables, hop = 2 (top 3)

city_record.hosting city(year, match id, host city)

county_public_safety.city(city id, county id, name, white, black, amerindian, asian, multiracial, hispanic)

world_1.country(code, name, continent, region, surface area, indepdent year, population, life expectancy, gnp, gnp old, local name, ...)

Spliced Questions

What is the most populace city that speaks English?; world_1.city(id, name, country code, district, population);

city_record.hosting city(year, match id, host city)

What is the most populace city that speaks English?; world_1.city(id, name, country code, district, population);

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

What is the most populace city that speaks English?; city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp);

city_record.hosting city(year, match id, host city)

Retrieved Tables, hop = 3 (top 3)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

city_record.hosting city(year, match id, host city)

world_1.city(id, name, country code, district, population)

MURRE ($r@3 = 100.0$)

Retrieved Tables, hop = 1 (top 3)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

world_1.city(id, name, country code, district, population)

e_government.addresses(address id, line 1 number building, town city, zip postcode, state province county, country)

Removal Questions

city_record.language(city id, language, percentage)

world_1.countrylanguage(countrycode, language, is official, percentage)

e_government.languages(language id, language name, language code, population)

Retrieved Tables, hop = 2 (top 3)

world_1.city(id, name, country code, district, population))

world_1.countrylanguage(countrycode, language, is official, percentage)

city_record.city(city id, city, hanzi, hanyu pinyin, regional population, gdp)

Removal Questions

None

None

None

(Early Stop)

Table 13: Detailed case study comparing MURRE with MURRE without Removal. The green means the relevant table, while the red means irrelevant. Each table is expressed in the form of “database name.tabel name(column names)”. r denotes recall.