

PromptMind Team at EHRSQL-2024: Improving Reliability of SQL Generation using Ensemble LLMs

Satya K Gundabathula
satyakesav123@gmail.com

Sriram R Kolar
sriramrakshithkolar@gmail.com

Abstract

This paper presents our approach to the EHRSQL-2024 shared task, which aims to develop a reliable Text-to-SQL system for electronic health records. We propose two approaches that leverage large language models (LLMs) for prompting and fine-tuning to generate EHRSQL queries. In both techniques, we concentrate on bridging the gap between the real-world knowledge on which LLMs are trained and the domain-specific knowledge required for the task. The paper provides the results of each approach individually, demonstrating that they achieve high execution accuracy. Additionally, we show that an ensemble approach further enhances generation reliability by reducing errors. This approach secured us 2nd place in the shared task competition. The methodologies outlined in this paper are designed to be transferable to domain-specific Text-to-SQL problems that emphasize both accuracy and reliability.

1 Introduction

Text-to-SQL technology translates natural language questions into executable SQL queries that can answer the questions using a provided database. A robust Text-to-SQL system could significantly increase productivity for anyone using databases by providing an easy-to-use natural language interface and reducing the need for expertise in different SQL dialects. These systems are particularly more valuable in domains where SQL knowledge is not essential, such as healthcare, where healthcare professionals like doctors, nurses, and hospital administrators spend a significant amount of time interacting with patient health records stored in databases.

In the era of Large Language Models (LLMs), the field of Text-to-SQL is gaining prominence as these models demonstrate impressive text generation capabilities without the need for fine-tuning.

Introduced in 2017, WikiSQL (Zhong et al., 2017) remains one of the largest datasets for Text-to-SQL and primarily caters to relatively simple queries. Subsequently, the SPIDER (Yu et al., 2018) and MULTI-SPIDER (Dou et al., 2023) datasets were developed. These datasets posed challenges with complex queries that required an understanding of the database schema and support for various languages. BIRD-Bench was introduced to bridge the gap between research and real-world applications by providing large and imperfect databases (Li et al., 2024). These datasets are good representations of typical Text-to-SQL tasks. However, the healthcare domain differs from these generic datasets for the following reasons:

- The questions asked by users maybe highly specialized and specific to the medical field.
- To answer such questions, systems must also possess an understanding of clinical terminology.
- Reliability is of paramount importance as errors can have serious consequences.

These differences present unique challenges for developing a reliable Text-to-SQL system for the healthcare domain. EHRSQL is the first dataset that closely captures the needs of hospital staff and serves appropriately for building and testing Text-to-SQL systems in the healthcare domain (Lee et al., 2022).

Our solution aims to create a Text to SQL system that emphasizes both reliability and accuracy. To achieve this, we divide the task into two phases:

- SQL Generation
- SQL Validation

In the first stage, we focus on SQL generation employing different techniques that include

prompting and fine-tuning of LLMs. In both approaches, we use the same prompting strategy to provide the LLM with database information and question-related context. Specifically, we use table schemas combined with sample column values as the database context, and similar questions from the training data as the task context. To identify similar questions from the training data, we employ an embedding-based similarity technique. Then, our goal is to maximize the LLM’s ability to generate highly accurate SQL statements utilizing this approach.

There are several reasons why LLMs may fail to generate correct SQL for a given question. Some common reasons include:

- Misinterpretation of question’s intent
- Incorrect assumptions or hallucinations about the database’s tables or columns
- Inaccuracies or hallucinations in the generated SQL query

Unlike many text generation tasks, Text-to-SQL tasks have a limited number of correct answers but potentially infinite incorrect ones. Inspired by this, we develop a second stage that evaluates the accuracy of the generated SQL. To evaluate the same, we propose an approach for Text-to-SQL that combines the results of multiple robust LLMs. Stronger LLMs often produce consistent outputs despite variations in temperature or other parameters, while smaller LLMs show lower consistency and accuracy. By leveraging the strengths of several robust LLMs, our approach minimizes the number of incorrect SQL queries and enhances the overall robustness and reliability of the Text-to-SQL system.

In the remainder of this paper, we discuss related work, introduce the EHRSQL-2024 task and dataset, and present our two-stage approach. We then provide the results of our experiments and conclude with a summary of our findings.

2 Related Work

Prior to the advent of LLMs, the primary focus of research in natural language processing involved refining specialized models using innovative strategies (Wang et al., 2020). Additionally, substantial efforts were devoted to developing sophisticated pre-training methodologies, such as those proposed by STAR (Cai et al., 2022), and exploring decoding strategies, as exemplified by PICARD (Scholak et al., 2021). However, these approaches typically

require substantial computational resources and novel techniques.

Large Language Models (LLMs) have been trained extensively on textual data, which has equipped them with vast knowledge. As a result, they exhibit exceptional probabilistic reasoning abilities and can excel at various tasks even without explicit training. Zero-shot prompting techniques, when used with LLMs, have not only narrowed the performance gap on Text-to-SQL but have also surpassed specialized pre-trained or fine-tuned models. Several prompt techniques have been developed based on this zero-shot approach for Text-to-SQL tasks, leading to remarkable achievements on datasets such as SPIDER (Dong et al., 2023), (Liu et al., 2023). Zero-shot generation capabilities can be further enhanced through techniques like in-context learning (ICL) and few-shot prompting.

DIN-SQL (Pourreza and Rafiei, 2023) adopts an in-context learning approach to break down complex SQL generation into manageable sub-tasks, leading to improved performance on intricate queries. Another technique, retrieval-augmented generation, provides relevant and helpful examples as a few-shot to guide SQL generation (Guo et al., 2024). These approaches have proven effective on general Text-to-SQL tasks but they have not yet been studied rigorously on domain-specific Text-to-SQL problems. Retrieval Augmented Fine-tuning (RAFT) introduces a novel fine-tuning technique that improves the in-domain performance of RAG while integrating domain-specific knowledge (Lewis et al., 2020).

Through our work, we delve into the application of these techniques for the EHRSQL-2024 task.

3 Shared Task and Dataset

The EHRSQL-2024 shared task (Lee et al., 2024) is aimed at creating a reliable SQL for answering questions posed in natural language on the MIMIC-IV demo database. The MIMIC-IV database consists of anonymized electronic health records of patients admitted to the Beth Israel Deaconess Medical Center. These records primarily cover two modules: hospital records and ICU records. The publicly available demo version of the database contains a subset of patient records for 100 individuals. It consists of 17 tables from both modules, encompassing a total of 109 columns.

	Total Samples	% Unanswerable
Train	5124	8.78 %
Valid	1163	19.95 %
Test	1167	19.97 %

Table 1: EHRSQL-2024 Dataset Statistics

3.1 Task Definition

The task aims to accurately generate SQL queries for answerable questions and predict null (ϕ) for unanswerable ones. Each correct answer earns a score of 1, while incorrect answers receive a score of $-c$, where c is the associated cost. The overall score RS for a cost c and prompt parameter θ can be expressed as below.

$$RS_{\theta}(C) = \sum_{i=1}^N \mathbb{1}(E(LLM_{\theta}(Q_i)) = E(GT_i)) - C * \mathbb{1}(E(LLM_{\theta}(Q_i)) \neq E(GT_i)) \quad (1)$$

where LLM represents the model that generates SQL based on a given question Q_i . GT_i denotes the ground truth SQL query for the question, and E signifies the executed value of the SQL query when run on a specific database. $\mathbb{1}$ is the indicator function.

The objective of this task is to find the optimal value of θ at a cost c with respect to the function $RS_{\theta}(C = c)$.

3.2 Dataset

The dataset contains a combination of answerable and unanswerable questions across three subsets: train, valid, and test. Table 1 provides an overview of the composition of each subset.

4 Approach

The reliable Text-to-SQL solution is decomposed into two stages as follows.

4.1 SQL Generation

To begin, we concentrate solely on boosting the number of accurately produced SQLs without being concerned with reducing the number of incorrect responses. As a result, the objective function becomes:

$$RS_{\theta}(C = 0) = \sum_{i=1}^N \mathbb{1}(E(LLM_{\theta}(Q_i)) = E(GT_i)) \quad (2)$$

Maximizing the success and minimizing hallucinations of the LLMs generation task require the

provision of the correct context. To achieve this, the following information is essential regarding the task at hand:

- **Database Schema** Comprising tables, columns, and their interrelationships, the database schema serves as a blueprint for the data stored in the database. This information guides the LLM in selecting the appropriate tables and columns.
- **Database Column Values** The actual values stored in the table columns offer additional information. This helps the LLM comprehend and perform operations such as data validation, manipulation, and filtering
- **Training Data** Providing questions (with corresponding SQL answers) similar to the current question aids the LLM in comprehending query formats, syntax, semantics, ambiguity resolution, and bridging the real-world knowledge gap with EHRSQL.

To produce SQL queries for each question, we employ an **in-context learning** approach. Here, the LLM is provided with similar question-SQL pairs, along with the relevant database content. To retrieve similar questions from the training data, we calculate cosine similarities between the evaluation question embedding and the training question embeddings.

We utilize the Angle model based on BERT, which aims to minimize the angle difference in a complex space (Li and Li, 2023). This approach helps overcome the negative impact caused by the saturation zone of the cosine function. The Angle embedding model ranks among the top 10 in the Massive Text Embedding Benchmark (MTEB), encompassing eight embedding tasks and 58 datasets (Muennighoff et al., 2022). While Angle effectively captures the semantic similarity between the intent of questions, it faces challenges in capturing the similarity between clinical terminology, which is also crucial for this task.

To bridge this gap, we combine Angle embeddings with PubMedBERT embeddings (Gu et al., 2020), trained on the PubMed literature. This allows us to enhance the system’s ability to capture clinical terminology. Since embedding similarity scores are not directly comparable across different models due to varying dimensionality, we perform z-normalization to ensure comparability. Algorithm 1 provides an overview of how we retrieve

the top N similar questions for a given question using two different embedding models.

To generate the SQL, we employed ICL and fine-tuning approaches with a consistent prompt template. A shorter version of the final prompt template is provided below for reference. For ICL, we utilized pre-trained models, such as GPT-4 (OpenAI et al., 2024) and Claude-3 Opus (Anthropic and others, 2024), with their default settings for temperature, top_p, and top_k parameters. To fine-tune GPT-3.5, we leveraged the retrieval augmented fine-tuning (RAFT) technique. For each training question, we generated similar questions using the multi-embedding retrieval approach while maintaining the prompt template. Given the limited size of the training set, we conducted fine-tuning with default parameters for only one epoch to prevent overfitting.

Prompt Template

This is a task converting a natural language question to an SQLite query for a database. You will be provided with the schema of the SQLite database followed by a few examples. You need to generate the SQLite query for a given question and you may return "null" if the question cannot be answered.

[Database Tables]

```
CREATE TABLE patients
(
  row_id int not null primary key, -- 42
  subject_id int not null unique, -- 201
  gender varchar(5) not null, -- 'm'
  dob timestamp(0) not null,
  dod timestamp(0)
);
```

...

[Examples]

[Q] : How many patients are there in total?

[SQL]: SELECT count(subject_id) FROM patients

[Q] : What is the gender of patient 1002?

[SQL]: SELECT gender FROM patients WHERE subject_id = 1002

[Q] : What is the date of birth of patient 1002?

SQL:

Figure 1 illustrates the complete process of generating SQL using an LLM post training for a given question.

4.2 SQL Validation

LLMs have a tendency to generate inaccurate and imaginary responses, regardless of the quality of the context they are provided. Therefore, we implement a second stage using an ensemble approach to eliminate errors generated during the initial generation stage. To verify whether the SQL generated by a two-model or three-model ensemble is correct, each query result is obtained by evaluating it against the database. Subsequently, the results are compared, and a match among all the results qualifies the query as correct.

5 Results

In this section, we present the comparison of the reliability scores of the individual models followed by ensembles.

5.1 Individual Models

Table 2 presents the reliability scores along with the percentage of unanswered questions for each model i.e. GPT-4, Claude-3 Opus and fine-tuned GPT-3.5.

Overall, Claude-3 Opus answered the most number of questions correctly while also answering them wrong more than others which led to the lowest RS10. GPT-4 appears to be more conservative in generating SQLs and has generated the most nulls. As refraining from generating for unanswerable questions is more important in this task, this led to achieving the best score on RS10 for GPT-4. Although the GPT-3.5 model is significantly less performant than GPT-4, the fine-tuned version brought the generation capability close to the GPT-4 model.

5.2 Ensemble

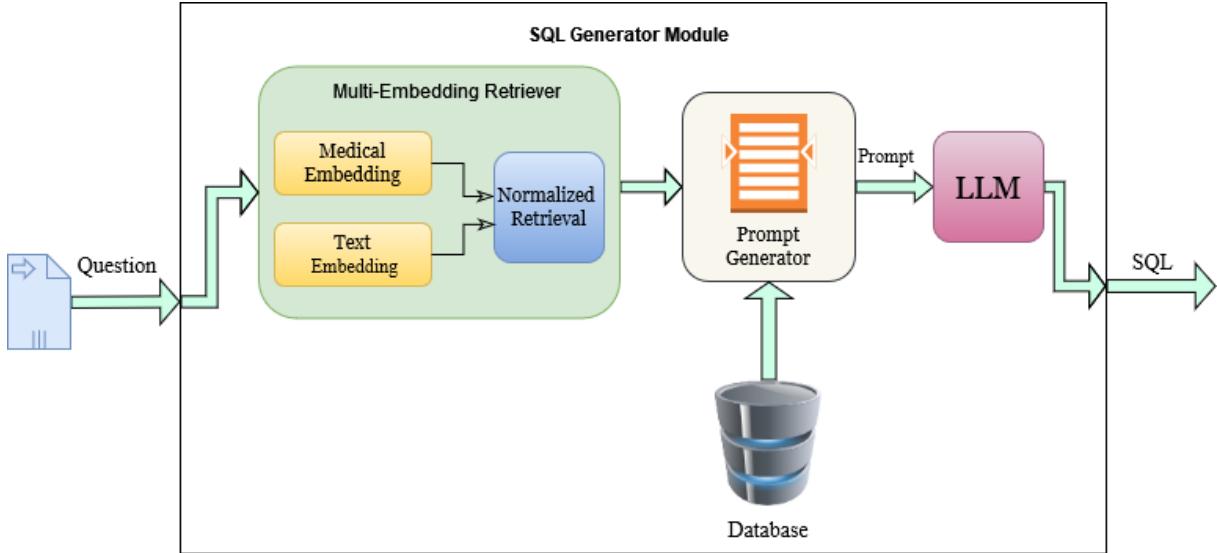
To select the ensemble model that achieves the best performance, we comprehensively evaluated all possible combinations of 2-model and 3-model ensembles. Table 3 provides a detailed comparison of the reliability scores achieved by these various model ensembles.

Among the 2-model ensembles, the combination of fine-tuned GPT-3.5 and Claude-3 Opus achieved the highest RS10 score, outperforming other models. Notably, the ensemble approach involving the fine-tuned GPT-3.5 model exhibited a significant

Algorithm 1: Multi-Embedding Retrieval

```
Data: train_questions, test_questions, N  
Result: Similar train questions for test questions  
// Same size as test_questions  
// Each element contains top N similar train questions and scores  
result  $\leftarrow$  [];  
foreach embed_model  $\in$  M do  
  train_embeddings  $\leftarrow$  create_embeddings(embed_model, train_questions);  
  test_embeddings  $\leftarrow$  create_embeddings(embed_model, test_questions);  
  questions, scores  $\leftarrow$  compute_similarity(test_embeddings, train_embeddings, top_n=N);  
   $\mu \leftarrow$  compute_mean(scores);  
   $\sigma \leftarrow$  compute_std(scores);  
  z_scores  $\leftarrow$  (scores  $- \mu$ ) /  $\sigma$ ;  
  // Sort and merge top N current questions with result  
  // If questions overlap, update with max score  
  result  $\leftarrow$  sort_and_merge(result, questions, z_scores);
```

Figure 1: SQL Generation Process



Model	RS0	RS10	Unanswered %
GPT-4	88.51	40.53	25.71
FT GPT-3.5	88.08	22.96	23.14
Opus	88.94	18.68	22.28

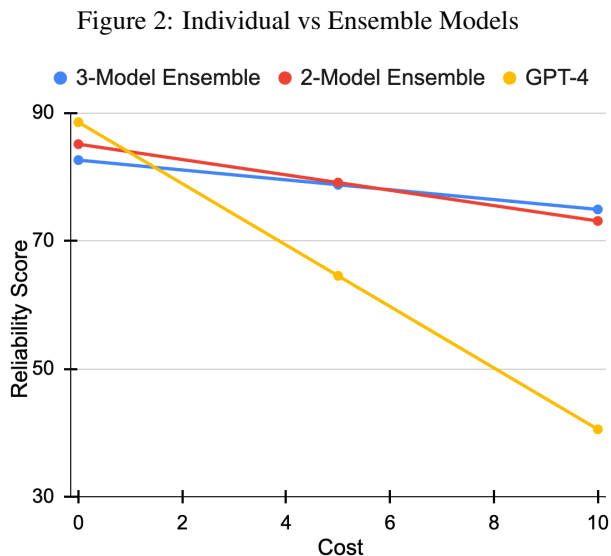
Table 2: Reliability Scores of Individual Models

reduction in errors compared to pre-trained models. This finding suggests that the fine-tuned model produces distinct errors from the pre-trained models, thus maximizing the validation benefits of ensemble approaches. The 3-model ensemble, however, achieved the best RS10 score among all approaches. To illustrate the effectiveness of Ensemble mod-

els, Figure 2 demonstrates the reliability scores of top-performing models from the individual, 2-model ensemble, and 3-model ensemble categories. When comparing against a stand-alone model, both 2-model and 3-model ensembles substantially minimize errors and obtain roughly equivalent but large RS10 scores. These results clearly demonstrate that ensemble approaches are effective validation mechanisms for creating reliable and accurate SQL generation systems.

Ensemble	RS0	RS10
GPT-4 + Opus	84.57	65.72
FT GPT-3.5 + GPT-4	84.83	71.97
FT GPT-3.5 + Opus	85.08	73.09
All	82.6	74.89

Table 3: Reliability Scores of Ensemble Models



6 Ablation Study

To assess the significance of each parameter in the final prompt employed for ICL and fine-tuning, we conduct an ablation study. In these experiments, we focus solely on pre-trained models because fine-tuning experiments are more expensive and time-consuming. To accelerate the process and maintain costs, we leverage GPT-3.5, a compact and less potent yet faster variant of the GPT family. Through these experiments, we extrapolate the efficacy of each parameter for prompting using more robust and advanced models such as GPT-4 and Claude-3 Opus. Table 4 provides the reliability scores obtained by progressively constructing a prompt with varying levels of complexity.

Incorporating few-shot examples in the prompt has substantially improved both the executable queries and reliability scores. This demonstrates the critical role of ICL with few-shot in Text-to-SQL tasks, particularly in the context of EHRSQL. The one-embedding few-shot experiment employs non-medical Angle embeddings (Li and Li, 2023), while the two-embeddings few-shot additionally leverages PubMedBERT (Gu et al., 2020). It is evident that adding medical embeddings enhances

Prompt Type	Executable %	RS0	RS10
No Few-shot	83.84	32.84	-440.06
One Embedding Few-Shot	95.89	66.98	7.65
Two Embeddings Few-Shot	98.34	69.13	11.52
Two Embeddings Few-Shot + Column Values	95.71	69.3	15.99

Table 4: Reliability Scores of GPT-3.5 with Different Prompt

all metrics by a good margin. While adding column values to the few-shot prompt decreased executable queries potentially leading to an increase in RS10, it also showed an improvement in RS0, indicating its usefulness as a signal. Through these experiments, we arrived at the final prompt, which enabled us to develop a highly reliable Text-to-SQL system.

7 Conclusion

Our work primarily aims to enhance the reliability of SQL generation, which is of paramount importance for the EHRSQL-2024 shared task. Although in-context learning with advanced LLMs such as GPT-4, Claude-3 Opus, or fine-tuning GPT-3.5 yields excellent RS0, errors still seem inevitable. The model’s ability to solve a specific task is heavily influenced by the training data. Repeatedly generating using the same prompt (or) the same model to validate often fails to minimize errors since hallucinations mainly originate from the training data. Fine-tuning GPT-3.5 resulted in different error tendencies compared to pre-trained models, even when using the same prompt. Therefore, ensemble LLMs, particularly those with a fine-tuned model, offer a superior approach for SQL validation, improving robustness and reliability. This approach has also secured us 2nd place in the competition.

8 Limitations and Risks

Our approach, while successful in this context, requires careful planning for real-world deployment due to certain limitations. Fine-tuning GPT-3.5 is computationally expensive and necessitates high-quality training data. Ensemble methods, though powerful for validation, introduce trade-offs in terms of cost and complexity. Crucially, it’s vital to evaluate potential biases inherited from the

LLM’s training data to ensure fair and reliable performance in practical applications.

References

- Anthropic and others. 2024. The claude 3 model family: Opus, sonnet, haiku. https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf. Online; accessed March 2024.
- Zefeng Cai, Xiangyu Li, Binyuan Hui, Min Yang, Bowen Li, Binhua Li, Zheng Cao, Weijie Li, Fei Huang, Luo Si, and Yongbin Li. 2022. **STAR: SQL guided pre-training for context-dependent text-to-SQL parsing**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1235–1247, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xuemei Dong, Chao Zhang, Yuhang Ge, Yuren Mao, Yunjun Gao, lu Chen, Jinshu Lin, and Dongfang Lou. 2023. **C3: Zero-shot text-to-sql with chatgpt**. *Preprint*, arXiv:2307.07306.
- Longxu Dou, Yan Gao, Mingyang Pan, Dingzirui Wang, Wanxiang Che, Dechen Zhan, and Jian-Guang Lou. 2023. **Multispider: Towards benchmarking multilingual text-to-sql semantic parsing**. In *AAAI Conference on Artificial Intelligence*.
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2020. **Domain-specific language model pretraining for biomedical natural language processing**.
- Chunxi Guo, Zhiliang Tian, Jintao Tang, Shasha Li, Zhihua Wen, Kaixuan Wang, and Ting Wang. 2024. **Retrieval-augmented gpt-3.5-based text-to-sql framework with sample-aware prompting and dynamic revision chain**. In *Neural Information Processing*, pages 341–356, Singapore. Springer Nature Singapore.
- Gyubok Lee, Hyeonji Hwang, Seongsu Bae, Yeonsu Kwon, Woncheol Shin, Seongjun Yang, Minjoon Seo, Jong-Yeup Kim, and Edward Choi. 2022. **Ehrsql: A practical text-to-sql benchmark for electronic health records**. *Advances in Neural Information Processing Systems*, 35:15589–15601.
- Gyubok Lee, Sunjun Kweon, Seongsu Bae, and Edward Choi. 2024. **Overview of the ehrsql 2024 shared task on reliable text-to-sql modeling on electronic health records**. In *Proceedings of the 6th Clinical Natural Language Processing Workshop*, Mexico City, Mexico. Association for Computational Linguistics.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. **Retrieval-augmented generation for knowledge-intensive NLP tasks**. *CoRR*, abs/2005.11401.
- Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, et al. 2024. **Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls**. *Advances in Neural Information Processing Systems*, 36.
- Xianming Li and Jing Li. 2023. **Angle-optimized text embeddings**. *arXiv preprint arXiv:2309.12871*.
- Aiwei Liu, Xuming Hu, Lijie Wen, and Philip S. Yu. 2023. **A comprehensive evaluation of chatgpt’s zero-shot text-to-sql capability**. *Preprint*, arXiv:2303.13547.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2022. **Mteb: Massive text embedding benchmark**. *arXiv preprint arXiv:2210.07316*.
- OpenAI, Josh Achiam, et al. 2024. **Gpt-4 technical report**. *Preprint*, arXiv:2303.08774.
- Mohammadreza Pourreza and Davood Rafiei. 2023. **Din-sql: Decomposed in-context learning of text-to-sql with self-correction**. *arXiv preprint arXiv:2304.11015*.
- Torsten Scholak, Nathan Schucher, and Dzmitry Bahdanau. 2021. **PICARD: Parsing incrementally for constrained auto-regressive decoding from language models**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9895–9901. Association for Computational Linguistics.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020. **RAT-SQL: Relation-aware schema encoding and linking for text-to-SQL parsers**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7567–7578, Online. 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*, Brussels, Belgium. Association for Computational Linguistics.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. **Seq2sql: Generating structured queries from natural language using reinforcement learning**. *CoRR*, abs/1709.00103.