
CHES: Contextual Harnessing for Efficient SQL Synthesis

Shayan Talaei
Stanford University
stalaei@stanford.edu

Mohammadreza Pourreza
University of Alberta
pourreza@ualberta.ca

Yu-Chen Chang
Stanford University
cjohnny@stanford.edu

Azalia Mirhoseini*
Stanford University
azalia@stanford.edu

Amin Saberi*
Stanford University
saberi@stanford.edu

Abstract

Utilizing large language models (LLMs) for transforming natural language questions into SQL queries (text-to-SQL) is a promising yet challenging approach, particularly when applied to real-world databases with complex and extensive schemas. In particular, effectively incorporating data catalogs and database values for SQL generation remains an obstacle, leading to suboptimal solutions. We address this problem by proposing a new pipeline that effectively retrieves relevant data and context, selects an efficient schema, and synthesizes correct and efficient SQL queries. To increase retrieval precision, our pipeline introduces a hierarchical retrieval method leveraging model-generated keywords, locality-sensitive hashing indexing, and vector databases. Additionally, we have developed an adaptive schema pruning technique that adjusts based on the complexity of the problem and the model’s context size. Our approach generalizes to both frontier proprietary models like GPT-4 and open-source models such as Llama-3-70B. Through a series of ablation studies, we demonstrate the effectiveness of each component of our pipeline and its impact on the end-to-end performance. Our method achieves new state-of-the-art performance on the cross-domain challenging BIRD dataset.

1 Introduction

Translating natural language questions into database queries, or text-to-SQL, is a long-standing research problem. This issue has been exacerbated in recent years due to the growing complexity of databases, driven by the increasing sizes of schemas (sets of columns and tables), values (content), and catalogs (metadata describing schemas and values) stored within them. Even the largest proprietary models, such as GPT-4, lag significantly behind human performance on text-to-SQL benchmarks, with a notable accuracy gap of 30% [Li et al., 2024b]. Beyond the complexity of writing SQL queries, this substantial gap is primarily caused by the need to effectively retrieve and integrate multiple sources of information, including database values, catalogs, and schema, each in different formats, which complicates the process.

In Figure 1, we show some of the challenges facing modern text-to-SQL systems. For instance, users’ questions might not directly match the stored values in the database, making it crucial to accurately identify the value format for effective SQL query formulation. Additionally, real-world database schemas often contain ambiguous column names, table names, and messy data, complicating the SQL translation process and necessitating a robust retrieval system to identify relevant information.

*Equal senior authorship

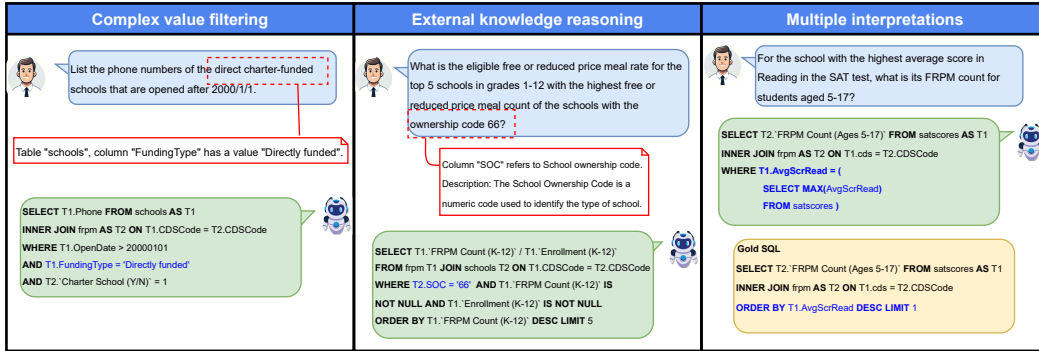


Figure 1: Example of challenges in text-to-SQL translation. 1) Questions passed by the users might not have the exact database value. 2) Column names might not be a good representation of what they store so using database catalogs is an essential part of text-to-SQL translation. 3) For a given question there are multiple ways of writing a correct SQL answer.

Moreover, there are typically multiple valid SQL queries that could answer the same question. For example, for the question illustrated on the right side of the Figure 1, one might use 'ORDER BY' and 'LIMIT 1' to find the highest average score, while another approach could involve a subquery with the 'MAX()' function, potentially leading to different outputs.

Earlier work in the area [Pourreza and Rafiei, 2024a, Wang et al., 2023, Qi et al., 2022, Rajkumar et al., 2022, Li et al., 2024c] has generally limited the context for SQL generation to table structures, column definitions, and sample rows. However, in production-level databases, the database catalog and database values constitute a rich source of information that is crucial for generating accurate SQL queries.

We introduce "CHESS: contextual harnessing for efficient SQL synthesis", an end-to-end text-to-SQL system that targets real-world and complex databases. CHESS introduces a scalable and effective LLM-based pipeline for SQL generation that consists of three main components: entity and context retrieval, schema selection, and SQL generation.

For entity and context retrieval, we present scalable and efficient methods using locality-sensitive hashing to retrieve database values from millions of rows, leveraging keyword detection, and vector databases to extract contextual information from database catalogs. Our approach utilizes both semantic and syntactic similarities between the database content and the user's query to enhance SQL prediction accuracy. In the schema selection phase, we utilize the retrieved information to narrow down the initial schema with potentially hundreds of columns to an efficient set of columns, usually less than ten. Throughout this step, we extract a minimal yet sufficient subset of the database schema. Finally, the extracted database schema is passed to a query generation module, which uses our fine-tuned SQL generator model combined with a revision step to effectively generate a SQL query.

Through a series of ablation studies 4.4, we demonstrate the critical role of each module in the pipeline in guiding LLMs to generate accurate SQL queries. Specifically, our entity and context retrieval module contributes substantially to performance, as evidenced by a 5% accuracy improvement.

At the time of submission, CHESS ranks first among all disclosed methodologies on BIRD [Li et al., 2024b], with a 65% and 66.69% execution accuracy on the development and test set respectively. To our knowledge, BIRD is the most challenging real-world text-to-SQL benchmark that is publicly available. It features more than 12000 unique question-SQL pairs, spanning 37 professional fields, including healthcare, education, blockchain, and sports, and covering 95 large databases with a combined size of 33.4 GB. An active leaderboard for BIRD is maintained by a third party (with a private test set), which we used to evaluate the performance of CHESS. CHESS also ranks second among all methods, with a marginal gap to the propriety (and undisclosed) approach that currently has a test set accuracy of 67.86% on BIRD's leaderboard.

We also provide an end-to-end open-source version of CHESS that obtains the best performance among other open-source baselines on the BIRD development set, with an execution accuracy of

61.5%. The use of open-source models for text-to-SQL is crucial, particularly when databases may contain private information that should not be shared with third-party LLM providers. All of the codes to reproduce the results reported in this paper are available on our Github repository².

Concretely, our contributions are as follows:

- Breaking down the Text-to-SQL task into a 3-staged pipeline, including entity and context retrieval, schema selection, and query generation
- A scalable hierarchical retrieval approach for extracting the important entities and contexts
- An efficient three-staged schema pruning protocol, consisting of individual column filtering, table selection, and a final column selection for extracting a minimally sufficient schema
- A fine-tuned open-source DeepSeek-33B Coder model with a novel training dataset construction approach with noise injection to mitigate error propagation
- A high-performing end-to-end open-source pipeline ensuring the privacy of the information
- Setting new state-of-the-art results on the BIRD dataset among known methodologies

2 Related Work

Generating accurate SQL queries from natural language questions, known as text-to-SQL is an active area of research within both the natural language processing (NLP) and database communities. Early efforts by the database community approach the problem through custom templates [Zelle and Mooney, 1996] marked initial advancements in this field, albeit at the expense of significant manual effort. Recently, text-to-SQL methodologies have increasingly incorporated transformer-based models, particularly sequence-to-sequence architectures [Vaswani et al., 2017, Sutskever et al., 2014]. These sequence-to-sequence models, capable of end-to-end training, are particularly well-suited for tasks requiring the generation of one sequence from another, such as translation, summarization, and text-to-SQL [Qin et al., 2022].

Initial sequence-to-sequence models, such as IRNet [Guo et al., 2019], utilized a bidirectional LSTM neural architecture to encode the query and employed self-attention to encode the database schema representation. To enhance the integration of schema information and capture its relationship with the question, models like RAT-SQL [Wang et al., 2019] and RASAT [Qi et al., 2022] have employed relation-aware self-attention mechanisms. Additionally, SADGA [Cai et al., 2021] and LGESQL [Cao et al., 2021] have adopted graph neural networks to represent the relational structures between the database schema and the queries. Although sequence-to-sequence models have improved performance, further advancements are necessary to bridge the gap with human performance. For example, none of the above techniques achieve an execution accuracy of more than 80% on the Spider hold-out test set [Yu et al., 2018].

Alongside the widespread adoption of LLMs across various NLP domains, the text-to-SQL field has similarly benefited from recent methodological innovations with LLMs to enhance performance. Early approaches [Rajkumar et al., 2022], leveraged the zero-shot in-context learning capabilities of LLMs for SQL generation. Building on this, subsequent models including DIN-SQL [Pourreza and Rafiei, 2024a], DAIL-SQL [Gao et al., 2023], MAC-SQL [Wang et al., 2023], and C3 [Dong et al., 2023] have enhanced LLM performance through task decomposition and techniques such as Chains of Thought (CoT) [Wei et al., 2022], self-consistency [Wang et al., 2022], and least-to-most prompting [Zhou et al., 2022]. In addition to in-context learning, proposals in DAIL-SQL [Gao et al., 2023], DTS-SQL [Pourreza and Rafiei, 2024b], and CodeS [Li et al., 2024a] have sought to elevate the capabilities of open-source LLMs through supervised fine-tuning, aiming to rival or exceed their larger, proprietary counterparts. However, the most notable performance gains have been observed in proprietary LLMs utilizing in-context learning methods [Li et al., 2024b]. Unlike previous efforts, this paper introduces a hybrid approach that combines both in-context learning and supervised fine-tuning to further enhance performance. Moreover, we propose novel methods to integrate contextual data such as database values and database catalog into the text-to-SQL pipeline, leveraging a rich yet often overlooked source of information.

Independently, but concurrent with our work, MCS-SQL [Lee et al., 2024] introduced a method that relies on using multiple prompts and sampling several responses from LLMs to mitigate sensitivity to

²<https://github.com/ShayanTalaie/CHES>

the arrangement of tables, columns, and few-shot in-context samples. They also devised a technique to filter out irrelevant table and column names. However, unlike our approach, they do not emphasize retrieving pertinent information from the database catalog and database values. Instead, they rely on using a large number of samples from LLMs to reduce prompt sensitivity for accuracy improvement.

3 Methodology

The text-to-SQL task is a challenging learning problem, requiring the model to translate a natural language question and database schema into a valid SQL query. The input consists of the text of the question, the database schema defining table structures and column types, and the database instance containing the actual or samples of the table content. Further, the input may include a database catalog with metadata, allowing free-form textual attribute values. Key challenges include noisy database content, the need to implicitly capture semantic correspondences between input elements, and the compositional nature of mapping language to SQL’s formal query structure.

"CHESS: Contextual Harnessing for Efficient SQL Synthesis" is an end-to-end text-to-SQL system designed for complex, real-world databases. It introduces a scalable and effective LLM-based pipeline consisting of three main components: entity and context retrieval, schema selection, and SQL generation. The entity and context retrieval component 3.1 uses keyword selection, locality-sensitive hashing, and vector databases to efficiently retrieve relevant database values and contextual information from large databases. The schema selection phase 3.2 then narrows down the initial schema to a minimal yet sufficient subset of columns. Finally, the extracted schema is passed to a query generation module 3.3 that utilizes a fine-tuned SQL generator model and a revision step to generate an accurate SQL query. The overall pipeline is demonstrated in Figure 2.

Finally, some of the implementation details of our method, including how we preprocess database values and data catalogs to expedite retrieval during the pipeline execution, are provided in Appendix C. Detailed execution traces of our pipeline are presented in Appendix E.

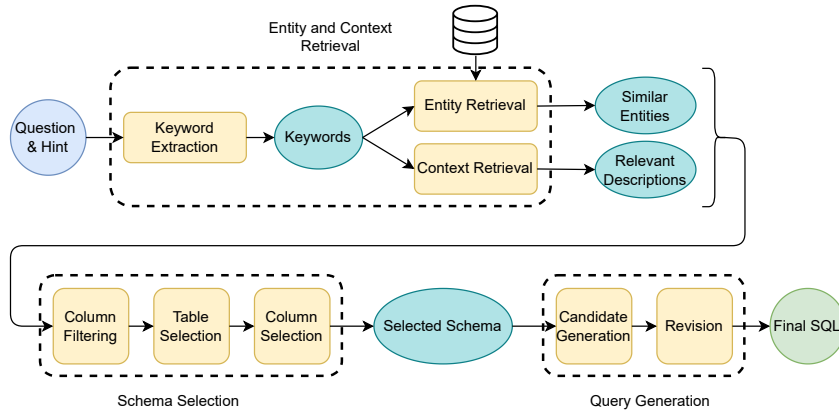


Figure 2: Our pipeline with modules for entity and context retrieval, schema selection, and query generation.

A key feature of our SQL query generation approach is the integration of multiple sources of information, including database catalogs, values, and schemas. However, these datasets with schemas containing hundreds of columns, each with detailed descriptions, and potentially millions of rows are often very large. Passing all this information to an LLM is often impractical due to limited context windows. Even if feasible, it can negatively impact the LLM’s reasoning capabilities, as demonstrated in [Hsieh et al., 2024] and our own ablation studies in Section 4.5. Our pipeline addresses this challenge by providing the LLM with minimal yet sufficient information necessary for each task. Maintaining *minimal sufficiency* is a key feature in all modules of the pipeline. Most crucially, during the SQL generation phase, we try to identify and pass to the model only the columns that are needed for the generation of the SQL query.

3.1 Entity and Context Retrieval

The first module in the pipeline identifies the relevant information in the input, including the entities referred to in the question and the contextual information provided about them in the database schema. This is done in three steps.

Keyword Extraction. To search for similar values in the database and schema descriptions, we first need to extract the main keywords from the natural-language question. Our approach to this problem is to prompt the model with few-shot examples of the task and the question, asking it to identify and extract keywords, keyphrases, and named entities.

Entity Retrieval. From the list of keywords extracted from the question, some may correspond to entities present in the database values. In this step, we search for similar values in the database and return the most relevant ones, along with their corresponding columns, for each keyword. As illustrated in Figure 1, searching for exact matches of the keywords cannot handle variations or typos, necessitating a more flexible search method. To measure the syntactic similarity between the keywords and the database values, we use the edit distance similarity metric. Additionally, to make the retrieval process more efficient, we propose a hierarchical retrieval strategy based on Locality Sensitive Hashing (LSH) and semantic (embedding) similarity measures, which we explain in detail in Appendix C. This approach allows us to efficiently retrieve values that exhibit a high degree of both syntactic and semantic similarity to the keywords.

Context Retrieval. In addition to the values, database catalogs explaining the schema may be available. For instance, each column may have a description, an extended column name (in the case of abbreviations), and a value description. As shown in Figure 1, this information can be useful, and not providing it to the model can lead to suboptimal performance. As explained before, in retrieving this context, we aim to identify only the minimally sufficient or the most relevant information. This is done by retrieving the most similar descriptions to the extracted keywords, measured by a semantic (embedding) similarity metric when querying the vector database of descriptions created during the preprocessing step.

3.2 Schema Selection

Our goal in this step is to narrow down the schema to include only the necessary tables and columns needed for generating the SQL query. We refer to this optimized set of necessary tables and columns as the *efficient schema*. Achieving an efficient schema leads to better performance in SQL query generation by excluding irrelevant information. We use recall and precision metrics to determine whether we have selected the correct tables and columns using the correct SQL query as the ground truth. The results are presented in Table 5 along with a detailed example of the schema selection process in Appendix F. The prompts for all of the sub-modules in schema selection are augmented with chain-of-thought prompting Wei et al. [2022] to improve the reasoning ability of LLMs which is essential for this task.

Individual Column Filtering. A database can contain hundreds of columns, many of which may be semantically irrelevant to the question. Looking for an efficient schema, we aim to filter out the irrelevant columns and pass only the most relevant ones to the table selection step. To accomplish this, we treat the relevance of each column to the question as a binary classification task for the model, essentially asking the LLM if the column may be relevant to the question. This step is useful only for removing the irrelevant columns that are obvious, but evaluating the relevance of a column in isolation is not always possible. We address this limitation in the subsequent steps, table selection and column selection, in which we give the model a more global view of the schema.

Table Selection. After filtering out irrelevant columns, we proceed to select the tables essential for generating the SQL query. In this step, we present the model with the filtered schema from the previous step and ask it to assess the relevance of each table, selecting only those necessary for the SQL query.

Final Column Selection. In the final step of schema selection, we aim to reduce the schema to the minimal set of columns necessary for generating the SQL query. We prompt the model to evaluate the necessity of each column in the filtered tables. This step includes a chain-of-thought explanation of why each column is needed, followed by the selection of the required columns.

3.3 Query Generation

At this point, we have selected an efficient schema augmented by the relevant context, containing all the necessary information to craft a SQL query that answers the question. In the following steps, we first write a candidate SQL query and then revise it to fix potential semantic and syntactic errors.

Candidate Generation. After reducing the schema to the minimal set of tables and columns, we prompt the model to generate an SQL query that answers the question. In the prompt, we provide the minimal schema obtained from the previous steps, along with the relevant values and descriptions retrieved in the first step of the pipeline. With this information, the model generates a candidate SQL query.

Revision. In the final step of the pipeline, we aim to fix potential logical and syntactic errors in the candidate SQL query. We provide the model with the database schema, the question, the generated candidate SQL query, and its execution result. The model is then asked to evaluate the correctness of the SQL query and revise it if necessary.

To assist the model in identifying and correcting mistakes, we give it a set of rules following [Bai et al., 2022]. In some cases, there may be multiple ways to correct the candidate SQL query. Even with zero-temperature sampling, the model might output different corrections across multiple samplings. To reduce the noise in the model’s output, we use self-consistency to select the SQL query that appears most consistently across three samples.

3.4 Preprocessing

To facilitate the information retrieval process outlined in 3.1 efficiently, we preprocess database values and catalogs before executing the pipeline. For database values, we conduct a syntactic search by creating a Locality Sensitive Hashing (LSH) index, as described in entity retrieval. For database catalogs, which contain longer texts requiring semantic understanding, we use a vector database retrieval method to measure semantic similarity.

Locality Sensitive Hashing Indexing of Values. To optimize the entity retrieval step, we employ a method capable of efficiently searching through large databases, which may contain millions of rows, to retrieve the most similar values. This step doesn’t require perfect accuracy but should retrieve a reasonably small set of similar values, such as a hundred elements. Locality Sensitive Hashing (LSH) is an effective technique for approximate nearest-neighbor searches. It allows us to retrieve database values that are most similar to a given keyword. During preprocessing, we index unique database values using LSH. Then, during the entity retrieval step, we query this index to quickly find the top similar values for a keyword.

Vector Database for Descriptions. As explained in context retrieval, extracting the most semantically relevant pieces of information from database catalogs is crucial for writing a SQL query. These documents can be extensive, with hundreds of pages explaining the entities and their relationships within the database, necessitating an efficient retrieval method. To perform a high-efficiency semantic similarity search, we preprocess the database catalogs into a vector database. During the context retrieval step, we query this vector database to find the most relevant pieces of information for the question at hand. For a more detailed description of our pipeline and the preprocessing phase, we encourage the reader to read Appendix C.

4 Experiments

4.1 Datasets and Metrics

The Spider dataset [Yu et al., 2018] includes 200 database schemas, with 160 schemas available for training and development, and 40 schemas reserved for testing. Notably, the databases used in the training, development, and test sets are distinct and do not overlap.

The recently introduced BIRD dataset [Li et al., 2024b] features 12,751 unique question-SQL pairs, covering 95 large databases with a combined size of 33.4 GB. This dataset spans 37 professional fields, including sectors such as blockchain, hockey, healthcare, and education. BIRD enhances SQL query generation by incorporating external knowledge and providing a detailed database catalog that includes column and database descriptions, thereby clarifying potential ambiguities. The SQL queries in BIRD are generally more complex than those found in the Spider dataset.

Subsampled Development Set (SDS). To facilitate ablation studies, reduce costs, and maintain the distribution of the BIRD development set, we subsampled 10% of each database in the development set, resulting in the Subsampled Development Set which we call SDS. This SDS consists of 147 samples: 81 simple, 54 moderate, and 12 challenging questions. For reproducibility, we included the SDS in our GitHub repository³.

Metrics. The Spider dataset evaluates SQL queries using two metrics: exact-set-match accuracy (EM) and execution accuracy (EX). Exact-set-match accuracy assesses each clause independently, requiring a perfect match with its corresponding clause in the reference SQL query. A SQL query is deemed correct only if it aligns completely with the reference across all components. However, this metric does not account for data values and has a high false negative rate due to multiple valid SQL formulations for a single question.

Execution accuracy (EX) evaluates the accuracy of the SQL output by comparing the results of the predicted query with those of the reference query when executed on specific database instances. This metric provides a more nuanced understanding of performance by accounting for variations in valid SQL queries that may arise from the same question.

Additionally, the BIRD benchmark introduces the Valid Efficiency Score (VES), which evaluates SQL query performance by considering both accuracy and execution speed. We achieve a similar ranking compared to other models on this metric as well but due to its high variance and dependence on the computational environment we exclude it from the current analysis.

4.2 BIRD Results

Since the test set of the BIRD benchmark is not available, we conducted our ablations and performance evaluations on the development set. We assessed our proposed method using both 1) proprietary and 2) open-source models. In the first scenario, we utilized our fine-tuned DeepSeek Coder model for candidate generation, GPT-3.5-turbo for column filtering, and GPT-4-turbo for the remaining LLM calls. We refer to this as our default engine set up. In the second scenario, we used our fine-tuned DeepSeek Coder model for candidate generation, with all other LLM calls handled by Llama-3-70B.

As reported in Table 1a, our approach using proprietary models achieved state-of-the-art execution accuracy on both the development and test sets of BIRD. Our method with open-source LLMs attained the highest performance among all open-source methods. At the time of the submission of this paper, the highest-performing method on the BIRD leaderboard is ExSL + granite-20b-code with an accuracy of 67.86% on the test set. Our approach ranks second with an accuracy of 66.69%.

Table 1: Performance of our proposed method on the BIRD development set and Spider test set, comparing to all published methods.

(a) BIRD results			(b) Spider test set	
Method	test EX	dev EX	Method	EX
CHESS + proprietary (ours)	66.69	65.00	MCS-SQL+GPT-4	89.6
MCS-SQL + GPT-4 [Lee et al., 2024]	65.45	63.36	CHESS (ours)	87.2
CHESS + Open LLMs (ours)	–	61.5	DAIL-SQL + GPT-4 [Gao et al., 2023]	86.6
SFT CodeS-15B [Li et al., 2024a]	60.37	58.47	DIN-SQL + GPT-4 [Pourreza and Rafiei, 2024a]	85.3
DTS-SQL + DeepSeek 7B [Pourreza and Rafiei, 2024b]	60.31	55.8	C3 + ChatGPT [Dong et al., 2023]	82.3
MAC-SQL + GPT-4 [Wang et al., 2023]	57.56	59.59	RESDSL-3B [Li et al., 2023]	79.9

³<https://github.com/ShayanTalaee/CHESS>

4.3 Spider Results

To evaluate the generalizability of our proposed method beyond the BIRD benchmark, we tested it on the Spider test set without specifically fine-tuning a new model for candidate generation or modifying the in-context learning samples. We followed our default engine setup. The only adjustment we made to our pipeline was the removal of the context retrieval node since the Spider test set lacks column or table descriptions, which are integral to our method. As shown in Table 1b, our approach achieved an execution accuracy of 87.2% on 2,147 samples from the test set, ranking it as the second-highest performing method among those published. This underscores the robustness of our method across different databases without any modifications. Notably, the best propriety (and undisclosed) method on the Spider test set leaderboard is Miniseek with an accuracy of 91.2%.

4.4 Ablation Studies

Models Ablation. Thanks to our efficient retrieval process, which carefully controls the number of tokens passed to LLMs, we can utilize an open-source LLM with a small context window size, specifically Llama-3 with only 8K tokens AI. This contrasts with previous works that predominantly use GPT-4 as their base model [Pourreza and Rafiei, 2024a, Lee et al., 2024, Wang et al., 2023]. In Table 2, we present the results of our proposed pipeline using various LLMs from different families on a subsampled development set dataset. The results indicate that our fine-tuned model for candidate generation significantly enhances performance. Notably, Llama-3’s performance surpasses that of GPT-3.5-turbo but does not yet reach the performance levels of GPT-4 in our analysis.

Table 2: This table shows the execution accuracy (EX) of different engine setups on the subsampled development set. Each engine setup is represented as a triplet (column filtering, candidate query generation, table/column selection + revision).

Engine setups	EX
(GPT-3.5-turbo, Fine-tuned DeepSeek, GPT-4-turbo)	64.62
(GPT-3.5-turbo, GPT-4-turbo, GPT-4-turbo)	55.78
(GPT-3.5-turbo, GPT-3.5-turbo, GPT-3.5-turbo)	49.65
(Llama-3-70B, Llama-3-70B, Llama-3-70B)	54.42
(Llama-3-70B, Fine-tuned DeepSeek, Llama-3-70B)	59.86

Table 3: The execution accuracy (EX) of the pipeline by removing each component on the (subsampled) dev set.

Pipeline Setup	EX	Δ EX
Full pipeline	64.62	–
w/o entity & context retrieval	59.86	-4.76
w/o individual column filtering	61.90	-2.72
w/o table selection	58.50	-6.12
w/o final column selection	59.18	-5.44
w/o revision	57.82	-6.80
with 1-time revision	61.22	-3.40

Modules Ablation. Table 3 presents the execution accuracy (EX), where different modules or components are omitted. In the configuration without entity and context retrieval, we retrieved a random example and included column descriptions for all columns. This approach highlights the significant impact of our selective retrieval, which outperforms naive context augmentation by 4.76% in execution accuracy. Additionally, we evaluated the effect of removing each submodule within the schema selection module, revealing that table selection is the most critical, contributing a 6.12% increase in performance. The table also illustrates the significant influence of the revision node, with a 6.80% improvement. Increasing the number of revision samples for self-consistency led to higher performance gains, aligning with findings from [Lee et al., 2024].

4.5 Performance Evaluation Across Queries with Varying Complexity

The BIRD benchmark categorizes questions and SQL pairs based on the number and type of SQL keywords used into three classes: easy, moderate, and challenging. In this section, we evaluate the performance of our method 1) using our fine-tuned model for candidate generation, and 2) using GPT-4 without our fine-tuned model. We compare the results to the original GPT-4 baseline as in the BIRD paper, where the question, evidence, and the complete schema with all tables and columns are presented to GPT-4 along with chain-of-thought reasoning prompts.

The analysis is conducted on the SDS dataset, and the results are detailed in Table 4. Our proposed method in both settings significantly improved performance across all classes. This analysis further demonstrates that providing an LLM with all available information can confuse the model and selective retrieval is crucial for achieving higher performance.

Table 4: Comparing the performance of our proposed method in two different settings with naively passing all information to GPT-4 across different difficulty levels on the subsampled development set.

	Easy	Moderate	Challenging	Overall
CHES (with fine-tuning)	65.43	64.81	58.33	64.62
CHES (w/o fine-tuning)	60.49	50.00	50.00	55.78
GPT-4-turbo (baseline)	54.32	35.18	41.66	46.25

4.6 Evaluation of the Schema Selection

Aside from the ablation studies to measure the effects of context retrieval, schema selection, and revision methods, we also measure their effect on the precision and recall of the tables and columns provided to the model for generating the final SQL. Precision and recall are calculated using the columns and tables that are used in the correct SQL queries as ground truth. As shown in Table 5, each step of the pipeline increases the precision of the selected tables and columns while only slightly decreasing the recall.

By achieving high recall and precision, the model has the most relevant information to generate the correct SQL in a small context window. An example of the reduction in the number of tables and columns of our pipeline is shown in Figure 12 where the final two tables are selected correctly and the final selected five columns include the two correct columns used by the correct SQL.

Table 5: Recall and Precision for Individual Column Filtering, Table Selection, and Column Selection compared to the tables and columns used in the correct SQL.

	Table		Column	
	Recall	Precision	Recall	Precision
No Filtering and Selection	1.0	0.33	1.0	0.11
Individual Column Filtering	1.0	0.33	0.98	0.21
Table Selection	0.97	0.89	0.96	0.45
Final Column Selection	0.96	0.90	0.94	0.71

5 Discussion and Limitations

In this paper, we propose a new LLM-powered pipeline, called CHES, for effective text-to-SQL synthesis. CHES consists of novel and efficient retrieval and schema pruning methodologies as well as scalable finetuning techniques. Our pipeline achieved state-of-the-art performance among all known methodologies on the challenging BIRD benchmark. Furthermore, we also developed an entirely open-source version of CHES, that for the first time, surpassed 60% execution accuracy on BIRD, narrowing the performance gap between closed-source and open-source LLMs, while ensuring the privacy of the data.

While our approach increases the capability of text-to-SQL methodologies, the ultimate goal is to bring full automation to the database querying process. For the challenging BIRD dataset, humans

still achieve a larger performance on query synthesis and future work should aim to further close this gap. We have quantified the efficacy of each of the components of our pipeline, which presents a guideline for strategies for further enhancements. Each of the information retrieval, schema pruning, and synthesis could benefit from further improvements, although as described in the paper, devising higher precision schema selection methodologies would be a high-impact area for future research that as it stands, might have an outsize impact on end-to-end accuracy.

References

- Meta AI. Exploring meta llama-3. <https://ai.meta.com/blog/meta-llama-3/>. Accessed: 2024-04-18.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- Ruichu Cai, Jinjie Yuan, Boyan Xu, and Zhifeng Hao. Sadga: Structure-aware dual graph aggregation network for text-to-sql. *Advances in Neural Information Processing Systems*, 34:7664–7676, 2021.
- Ruisheng Cao, Lu Chen, Zhi Chen, Yanbin Zhao, Su Zhu, and Kai Yu. Lgesql: line graph enhanced text-to-sql model with mixed local and non-local relations. *arXiv preprint arXiv:2106.01093*, 2021.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
- Xuemei Dong, Chao Zhang, Yuhang Ge, Yuren Mao, Yunjun Gao, Jinshu Lin, Dongfang Lou, et al. C3: Zero-shot text-to-sql with chatgpt. *arXiv preprint arXiv:2307.07306*, 2023.
- Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. Text-to-sql empowered by large language models: A benchmark evaluation. *arXiv preprint arXiv:2308.15363*, 2023.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y Wu, YK Li, et al. Deepseek-coder: When the large language model meets programming—the rise of code intelligence. *arXiv preprint arXiv:2401.14196*, 2024.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. Towards complex text-to-sql in cross-domain database with intermediate representation. *arXiv preprint arXiv:1905.08205*, 2019.
- Cheng-Ping Hsieh, Simeng Sun, Samuel Kriman, Shantanu Acharya, Dima Rekeshe, Fei Jia, and Boris Ginsburg. Ruler: What’s the real context size of your long-context language models? *arXiv e-prints*, pages arXiv–2404, 2024.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Dongjun Lee, Choongwon Park, Jaehyuk Kim, and Heesoo Park. Mcs-sql: Leveraging multiple prompts and multiple-choice selection for text-to-sql generation. *arXiv preprint arXiv:2405.07467*, 2024.
- Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. Resdsq: Decoupling schema linking and skeleton parsing for text-to-sql. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13067–13075, 2023.
- Haoyang Li, Jing Zhang, Hanbing Liu, Ju Fan, Xiaokang Zhang, Jun Zhu, Renjie Wei, Hongyan Pan, Cuiping Li, and Hong Chen. Codes: Towards building open-source language models for text-to-sql. *arXiv preprint arXiv:2402.16347*, 2024a.
- Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, et al. 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, 2024b.
- Zhishuai Li, Xiang Wang, Jingjing Zhao, Sun Yang, Guoqing Du, Xiaoru Hu, Bin Zhang, Yuxiao Ye, Ziyue Li, Rui Zhao, et al. Pet-sql: A prompt-enhanced two-stage text-to-sql framework with cross-consistency. *arXiv preprint arXiv:2403.09732*, 2024c.
- OpenAI. Embeddings, n.d. Retrieved May 15, 2024, from <https://platform.openai.com/docs/guides/embeddings>.

- Mohammadreza Pourreza and Davood Rafiei. Din-sql: Decomposed in-context learning of text-to-sql with self-correction. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Mohammadreza Pourreza and Davood Rafiei. Dts-sql: Decomposed text-to-sql with small large language models. *arXiv preprint arXiv:2402.01117*, 2024b.
- Jiexing Qi, Jingyao Tang, Ziwei He, Xiangpeng Wan, Yu Cheng, Chenghu Zhou, Xinbing Wang, Quanshi Zhang, and Zhouhan Lin. Rasat: Integrating relational structures into pretrained seq2seq model for text-to-sql. *arXiv preprint arXiv:2205.06983*, 2022.
- Bowen Qin, Binyuan Hui, Lihan Wang, Min Yang, Jinyang Li, Binhua Li, Ruiying Geng, Rongyu Cao, Jian Sun, Luo Si, et al. A survey on text-to-sql parsing: Concepts, methods, and future directions. *arXiv preprint arXiv:2208.13629*, 2022.
- Nitarshan Rajkumar, Raymond Li, and Dzmitry Bahdanau. Evaluating the text-to-sql capabilities of large language models. *arXiv preprint arXiv:2204.00498*, 2022.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. Rat-sql: Relation-aware schema encoding and linking for text-to-sql parsers. *arXiv preprint arXiv:1911.04942*, 2019.
- Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, Qian-Wen Zhang, Zhao Yan, and Zhoujun Li. Mac-sql: Multi-agent collaboration for text-to-sql. *arXiv preprint arXiv:2312.11242*, 2023.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. *arXiv preprint arXiv:1809.08887*, 2018.
- John M Zelle and Raymond J Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the national conference on artificial intelligence*, pages 1050–1055, 1996.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022.

A Query Generation with the Correct Context

We also fine-tuned the query generator module with the correct contextual information including only the relevant columns and tables and their descriptions. As shown in 6, the performance with the correct contextual information based on the gold SQL reached 72.4%, underscoring the critical role of retrieving the efficient schema information.

Table 6: This table shows the maximum execution accuracy (EX) possible for our candidate SQL module generation by passing it the correct context for questions in the BIRD dataset.

Engine	EX
CHESS	64.62
CHESS + correct context	72.4

B Finetuning the SQL Generator

B.1 Fine-tuning Dataset and Model

To enhance the generation of better candidate SQL queries prior to our revision module, we fine-tuned the DeepSeek Coder 34B Guo et al. [2024] on the training set of the BIRD benchmark, which comprises approximately 9,500 samples. In constructing the fine-tuning dataset, rather than solely using correct tables and columns like the work proposed in Pourreza and Rafiei [2024b], we developed a heuristic to address the error propagation issue. Recognizing that previous steps in the pipeline may not always pinpoint the most efficient schema with perfect accuracy, we intentionally introduced some noise into our dataset creation to train the model. Specifically, we included columns and tables that were incorrect but shared similar naming conventions and semantic attributes with the correct schema. We also utilized our keyword selection module to extract keywords from the questions, search for these keywords in the database, and incorporate them into the prompt.

B.2 Hyperparameters

We fine-tuned the candidate SQL generator model using 4-bit quantization of the base model and LORA adapters Hu et al. [2021], a technique formally referred to as QLORA Dettmers et al. [2024]. We configured the LORA rank parameter to 128 and set the LORA alpha parameter to 256. The fine-tuning process was conducted over two epochs on the constructed dataset, utilizing a batch size of 32 and a learning rate of $1e-4$, along with a cosine scheduler, all on a single H100 GPU for 4 hours.

C Implementation Details

C.1 Locality Sensitive Hashing Indexing of Database Values

Our goal in the entity retrieval sub-module is to retrieve database values that most closely match a set of keywords derived from the question. It is important to recognize that keywords from the question may not exactly correspond to the database values due to potential typos, variations in expression, or the common scenario where users are unaware of the precise format used to store data in the database. This reality demands a retrieval strategy that is both robust and adaptable, capable of accommodating such discrepancies. Relying solely on exact match retrieval, as suggested in prior studies [Li et al., 2024a], may not be sufficiently effective.

To address this, we employ string similarity measures, such as edit distance and semantic embedding, to retrieve the values most similar to the keywords. However, computing the edit distance and embedding similarity for every keyword against all values in the database is computationally expensive and time-consuming. To balance efficiency and accuracy, we utilize a hierarchical retrieval method.

Locality Sensitive Hashing (LSH) is an efficient technique for approximate nearest neighbor searches, which allows us to retrieve the most similar values to a keyword in the database. In the pre-processing stage, we index unique values in the database using LSH. Then, in the entity retrieval step of our

pipeline, we query this index to rapidly find the top similar values to a keyword. Our approach involves using LSH queries to retrieve the top 10 similar values, after which we compute the edit distance and semantic similarity between the keyword and these values to further refine the results.

To simultaneously utilize edit distance and embedding similarity, we first identify the top 10 values closest to each keyword based on cosine similarity between their embedding vectors (obtained using OpenAI text-embedding-3-small OpenAI [n.d.]) and the keyword’s embedding vector. We then filter out values that fall below a specific threshold. Finally, for each keyword and column, we retain only the value that has the smallest edit distance.

We observed a significant reduction in time complexity, from 5 minutes to 5 seconds, using this method compared to a naive approach of computing the edit distance for all unique values in the database on the fly. While computing edit distance is proportional to the size of the database—significantly increasing the time complexity for processing a single question—using LSH allows us to index values in the pre-processing step and, during entity retrieval, rapidly query the index to find the most similar values to a keyword in a much more time-efficient manner.

C.2 Vector database

Each database schema in the BIRD benchmark Li et al. [2024b] includes detailed descriptions for columns, specifying the contents of each column and the values for categorical columns. Providing these descriptions to the model is essential for guiding the SQL query generation process. However, incorporating all descriptions in the prompt can overwhelm the model, potentially leading to the generation of incorrect SQL queries, as observed in section 4.5. It is important to note that the database catalog in the BIRD benchmark provides a relatively limited view of database metadata. In contrast, real-world production-level databases often contain more diverse information, including value ranges, constraints, and usage instructions for each table. Our proposed method can effectively utilize this extensive metadata to enhance performance.

To evaluate the relevance of descriptions to a given question, we employ embedding similarity [OpenAI, n.d.], which quantifies the semantic similarity between the question and each description. To enhance the efficiency of the retrieval process, we pre-process the descriptions and created the embedding vectors for each of them and stored in a vector database, utilizing ChromaDB in our implementation. During the context retrieval phase of our pipeline, we query this vector database to identify descriptions that are most semantically aligned with the question. This targeted approach ensures that only the most pertinent information is provided to the model, thereby improving the accuracy of the generated SQL queries.

C.3 Local Column filtering

In the local column filtering module, decisions to retain a column for subsequent steps are made independently, without considering relative information. The data provided to the LLM for column filtering includes: 1) the table name, 2) the column name, 3) the data type, 4) descriptions, if retrieved during the context retrieval step, and 5) database values, if retrieved during the entity retrieval step. To enhance the model performance and make the task definition clear to the model, we used few-shot samples for this sub-module. This method is similar to the proposed approach in [Li et al., 2024a] model. The difference in our approach is that we complimented the local column filtering with column selection which can further reduce the number of columns by taking the relative information into account as well.

Some key columns for SQL generation, which we call linking columns, such as those with foreign and primary key constraints, are crucial for writing SQL queries. For instance, questions about counting entities often requires primary keys, and joining tables necessitates foreign key columns. However, in the column filtering and selection sub-modules, some of these essential columns may be initially rejected because they do not semantically relate to the given question. Despite this, these columns are indispensable for SQL generation. Therefore, in all of our sub-modules, we consistently retain foreign key and primary key columns, irrespective of the outputs from column selection and filtering processes.

C.4 Revision

Revising generated candidate SQL queries is a critical aspect of our proposed pipeline. In addition to the database schema, the question, and the candidate SQL query, we also provide the execution result of the SQL query. This gives the LLM an opportunity to view the retrieved data and revise the SQL query accordingly. This process mirrors human behavior when writing complex SQL queries; typically, we start with a draft query and refine it based on the results of its execution. Furthermore, this method allows the LLM to make necessary adjustments to the SQL query in instances of execution syntax errors.

In this step, we also incorporate instructions derived from our error analysis G to guide the model towards generating correct SQL queries. For instance, as shown in 3, we guided the model to ensure that all requested columns are included in the SQL query. In this specific example, the revision model identified a missing column and successfully added it to the query.

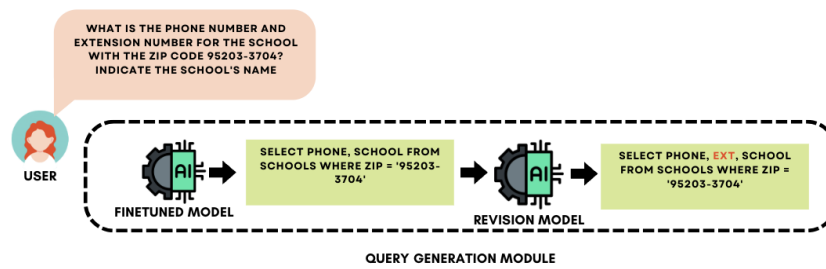


Figure 3: An example of using instructions to guide the revision model to fix missing columns in the candidate query.

D Prompt Templates

In this section we provide the exact prompts that have been used for each of the sub-modules in our pipeline. For easier parsing of the LLM's output, we instruct the model to generate a valid JSON object in its response if needed.

```
Objective: Analyze the given question and hint to identify and extract
keywords, keyphrases, and named entities. These elements are crucial
for understanding the core components of the inquiry and the guidance
provided. This process involves recognizing and isolating significant
terms and phrases that could be instrumental in formulating searches or
queries related to the posed question.

Instructions:
1. Read the Question Carefully: Understand the primary focus and
specific details of the question. Look for any named entities (such as
organizations, locations, etc.), technical terms, and other phrases that
encapsulate important aspects of the inquiry.
2. Analyze the Hint: The hint is designed to direct attention toward
certain elements relevant to answering the question. Extract any
keywords, phrases, or named entities that could provide further clarity
or direction in formulating an answer.
3. List Keyphrases and Entities: Combine your findings from both
the question and the hint into a single Python list. This list should
contain:
- Keywords: Single words that capture essential aspects of the question
or hint.
- Keyphrases: Short phrases or named entities that represent specific
concepts, locations, organizations, or other significant details.
Ensure to maintain the original phrasing or terminology used in the
question and hint.

{FEWSHOT_EXAMPLES}

Task:
Given the following question and hint, identify and list all relevant
keywords, keyphrases, and named entities.

Question: {QUESTION}

Hint: {HINT}

Please provide your findings as a Python list, capturing the essence
of both the question and hint through the identified terms and phrases.
Only output the Python list, no explanations needed.
```

Figure 4: Template for Keyword and Entity Extraction

You are an expert and very smart data analyst.
Your task is to analyze the provided database schema, comprehend the posed question, and leverage the hint to identify which tables are needed to generate a SQL query for answering the question.

Database Schema Overview:
{DATABASE_SCHEMA}

This schema provides a detailed definition of the database's structure, including tables, their columns, primary keys, foreign keys, and any relevant details about relationships or constraints.
For key phrases mentioned in the question, we have provided the most similar values within the columns denoted by "- examples" in front of the corresponding column names. This is a critical hint to identify the tables that will be used in the SQL query.

Question:
{QUESTION}

Hint:
{HINT}

The hint aims to direct your focus towards the specific elements of the database schema that are crucial for answering the question effectively.

Task:
Based on the database schema, question, and hint provided, your task is to determine the tables that should be used in the SQL query formulation. For each of the selected tables, explain why exactly it is necessary for answering the question. Your explanation should be logical and concise, demonstrating a clear understanding of the database schema, the question, and the hint.

Please respond with a JSON object structured as follows:

```
{
  "chain_of_thought_reasoning": "Explanation of the logical analysis
that led to the selection of the tables.",
  "table_names": ["Table1", "Table2", "Table3", ...]
}
```

Note that you should choose all and only the tables that are necessary to write a SQL query that answers the question effectively.
Take a deep breath and think logically. If you do the task correctly, I will give you 1 million dollars.

Only output a json as your response.

Figure 5: Template for the table selection module.

You are an expert and very smart data analyst.
Your task is to examine the provided database schema, understand the posed question, and use the hint to pinpoint the specific columns within tables that are essential for crafting a SQL query to answer the question.

Database Schema Overview:
{DATABASE_SCHEMA}

This schema offers an in-depth description of the database's architecture, detailing tables, columns, primary keys, foreign keys, and any pertinent information regarding relationships or constraints. Special attention should be given to the examples listed beside each column, as they directly hint at which columns are relevant to our query. For key phrases mentioned in the question, we have provided the most similar values within the columns denoted by "- examples" in front of the corresponding column names. This is a critical hint to identify the columns that will be used in the SQL query.

Question:
{QUESTION}

Hint:
{HINT}

The hint aims to direct your focus towards the specific elements of the database schema that are crucial for answering the question effectively.

Task:
Based on the database schema, question, and hint provided, your task is to identify all and only the columns that are essential for crafting a SQL query to answer the question. For each of the selected columns, explain why exactly it is necessary for answering the question. Your reasoning should be concise and clear, demonstrating a logical connection between the columns and the question asked.

Tip: If you are choosing a column for filtering a value within that column, make sure that column has the value as an example.

Please respond with a JSON object structured as follows:

```
{
  "chain_of_thought_reasoning": "Your reasoning for selecting the columns,
  be concise and clear.",
  "table_name1": ["column1", "column2", ...],
  "table_name2": ["column1", "column2", ...],
  ...
}
```

Make sure your response includes the table names as keys, each associated with a list of column names that are necessary for writing a SQL query to answer the question.

For each aspect of the question, provide a clear and concise explanation of your reasoning behind selecting the columns.

Take a deep breath and think logically. If you do the task correctly, I will give you 1 million dollars.

Only output a json as your response.

Figure 6: Template for the column selection module.

You are a detail-oriented data scientist tasked with evaluating the relevance of database column information for answering specific SQL query question based on provided hint.

Your goal is to assess whether the given column details are pertinent to constructing an SQL query to address the question informed by the hint. Label the column information as "relevant" if it aids in query formulation, or "irrelevant" if it does not.

Procedure:

1. Carefully examine the provided column details.
2. Understand the question about the database and its associated hint.
3. Decide if the column details are necessary for the SQL query based on your analysis.

Here is an example of how to determine if the column information is relevant or irrelevant to the question and the hint:

{FEWSHOT_EXAMPLES}

Now, it's your turn to determine whether the provided column information can help formulate a SQL query to answer the given question, based on the provided hint.

The following guidelines are VERY IMPORTANT to follow. Make sure to check each of them carefully before making your decision:

1. You're given only one column's information, which alone isn't enough to answer the full query. Concentrate solely on this provided data and assess its relevance to the question and hint without considering any missing information.
2. Read the column information carefully and understand the description of it, then see if the question or the hint is asking or referring to the same information. If yes then the column information is relevant, otherwise it is irrelevant.

...

Column information:

{COLUMN_PROFILE}

Question:

{QUESTION}

HINT:

{HINT}

Take a deep breath and provide your answer in the following json format:

```
{  "chain_of_thought_reasoning": "One line explanation of why or why not the column information is relevant to the question and the hint.",  "is_column_information_relevant": "Yes" or "No"}
```

Only output a json as your response.

Figure 7: Template for the column filtering module.

You are a data science expert.
Below, you are presented with a database schema and a question.
Your task is to read the schema, understand the question, and generate a valid SQLite query to answer the question.
Before generating the final SQL query think step by step on how to write the query.

Database Schema:
{DATABASE_SCHEMA}

This schema offers an in-depth description of the database's architecture, detailing tables, columns, primary keys, foreign keys, and any pertinent information regarding relationships or constraints. Special attention should be given to the examples listed beside each column, as they directly hint at which columns are relevant to our query.

Database admin instructions:
{DATABASE_ADMIN_INSTRUCTIONS}

Question:
{QUESTION}

Hint:
{HINT}

Please respond with a JSON object structured as follows:

```
{
  "chain_of_thought_reasoning": "Your thought process on how you arrived at the final SQL query.",
  "SQL": "Your SQL query in a single string."
}
```

Priority should be given to columns that have been explicitly matched with examples relevant to the question's context.

Take a deep breath and think step by step to find the correct SQLite SQL query. If you follow all the instructions and generate the correct query, I will give you 1 million dollars.

Figure 8: Template for SQL Query Candidate Generation

```
Objective: Your objective is to make sure a query follows the database
admin instructions and use the correct conditions.

Database Schema:
{DATABASE_SCHEMA}

Database admin instructions:
{DATABASE_ADMIN_INSTRUCTIONS}

{MISSING_ENTITIES}

Question:
{QUESTION}

Hint:
{EVIDENCE}

Predicted query:
{SQL}

Query result:
{QUERY_RESULT}

Please respond with a JSON object structured as follows (if the sql query
is correct, return the query as it is):

{
  "chain_of_thought_reasoning": "Your thought process on how you arrived
at the solution. You don't need to explain the instructions that are
satisfied.",
  "revised_SQL": "Your revised SQL query."
}

Take a deep breath and think step by step to find the correct SQLite SQL
query. If you follow all the instructions and generate the correct query,
I will give you 1 million dollars.
```

Figure 9: Template for SQL Query revision

E Execution Trace

In this section, we provide detailed execution traces of our pipeline, showcasing the output of each module. For example, as shown in figure 10, the effectiveness of our keyword extraction module is highlighted. It successfully identifies crucial values, such as "Lewis Hamilton," which are necessary for constructing SQL queries. These identified keywords aid in retrieving database values and relevant column descriptions for input to the LLM.

In the entity retrieval sub-module, each keyword extracted is first broken down into individual words by splitting on spaces. This process makes our search order invariant, as entities mentioned in the question do not necessarily follow the same format as those in the database. Subsequently, for each keyword, the most similar database values are extracted. Likewise, all detected keywords are used to filter column descriptions, which are then provided for subsequent steps. This strategic provision of relevant information facilitates further processes, such as schema selection and query generation, thereby enhancing the accuracy of schema detection and SQL query formulation.

In 11, we present a more challenging scenario where the question and evidence do not directly reference the relevant column names. The query mentions "schools with the ownership code 66," yet no column explicitly includes "ownership" in its name. In this case, the relevant column is "SOC" column from the schools table. The connection between this column and the question can only be discerned through the semantic similarity between the column descriptions and the question. Here, our context retrieval node plays a crucial role, as it effectively retrieves the relevant information from the database catalog, which is essential for answering this question.

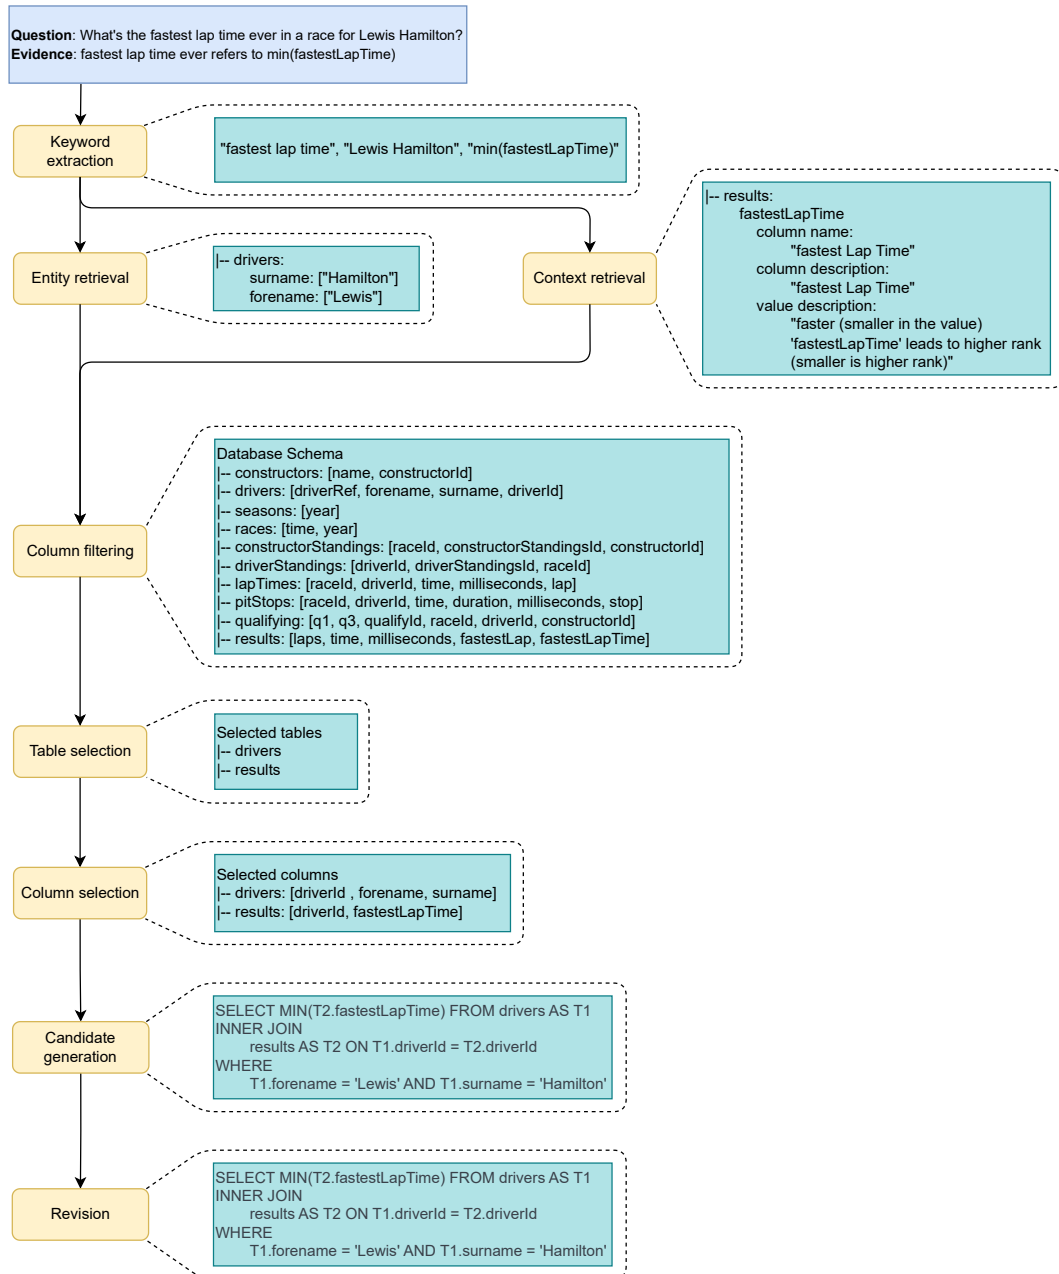


Figure 10: Execution flowchart for instance formula_1_926

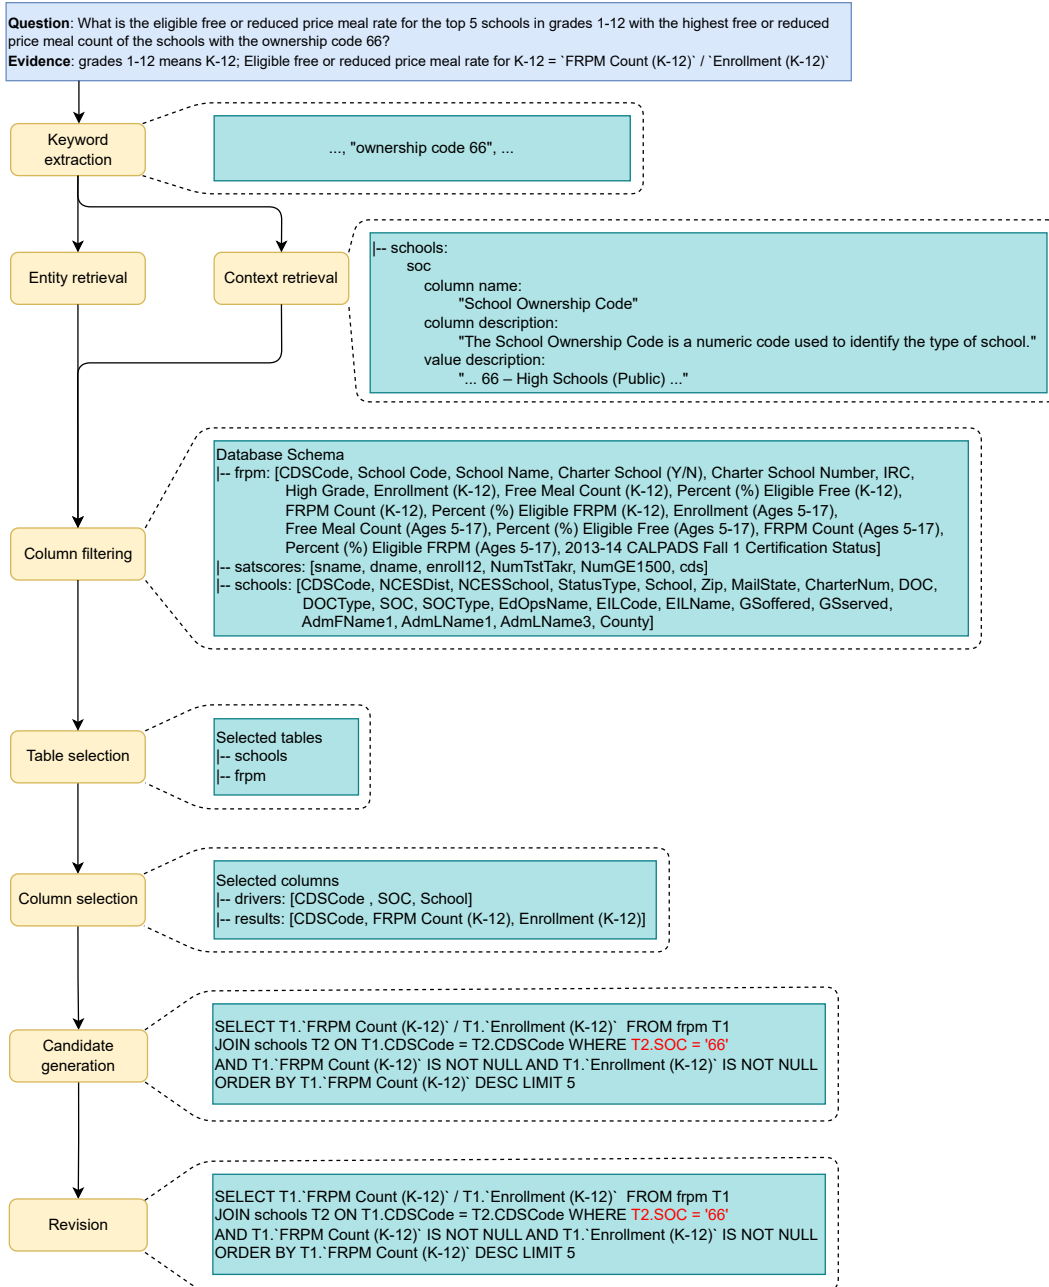


Figure 11: Execution flowchart for instance California Schools_32

F Schema Selection Example

As described in the methodology 3, following the *minimal sufficiency rule* we devoted the second phase of our pipeline to schema selection 3.2. In this part, we use an example to showcase how we narrow down the initial schema throughout the local column filtering, the table selection, and the column selection.

The example is chosen from *Formula 1* database from the BIRD benchmark, with a total of 13 tables and 96 columns. Here is the question and its evidence:

- **Question:** What's the fastest lap time ever in a race for Lewis Hamilton?
- **Evidence:** fastest lap time ever refers to `min(fastestLapTime)`

Figure 12 shows the number of tables and columns that are considered as the sub-selected schema after each step. Starting from 13 tables and 96 columns, these numbers reduced to 36 columns in 13 tables after the column filtering step. Subsequently, table selection narrowed it down further to 2 tables and 7 columns. Finally, column selection yielded the final schema with 2 tables and 5 columns, which is used for the SQL generation.

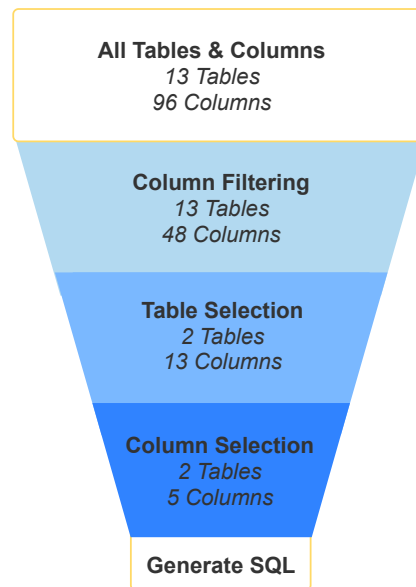


Figure 12: Funnel graph illustrating the progressive narrowing down of database schema through column filtering, table selection, and column selection steps, leading to the final schema used for SQL generation.

To further illustrate the details of the schema selection process, we use the entity-relationship diagram (ERD) 13. In this figure, each table is represented as a block with its columns listed below it. Primary keys are underlined and the foreign keys are in italics connecting the corresponding columns. As shown in the legend, the columns that remained in the selected schema after each step are colored with a gradient of white to dark blue; white represents the columns that were present in the initial schema and got filtered after column filtering while dark blue shows the columns that are selected after column selection to be passed for SQL candidate generation.

There are some points worth emphasizing in the plot. First, the decision to filter linking columns (Primary and foreign keys) is a task requiring a global view of the schema and cannot be done in the local view of column filtering, hence we do not filter these columns and include all of them in the result of column filtering step, explaining why all of the primary and foreign keys are present after this step. Second, some columns such as "laps", "time", and "milliseconds" are all semantically related to the question because question asks for fastest lap time, which shows the column filtering module successfully use the local information to find all relevant columns. However, all of these columns are not going to be used for crafting the SQL query, so we need to find the relevant columns with respect to their relative information, which is going to be done in the table selection and column

selection modules. In the table selection step, as it can be observed from the figure, the “lapTimes” table which has all information about time has been dropped by table selection since there is a more relevant column, “fastestLapTime”, which can be used to answer the question. This was a concrete example, which shows how local and global view of the columns and tables can help to pinpoint the correct schema.

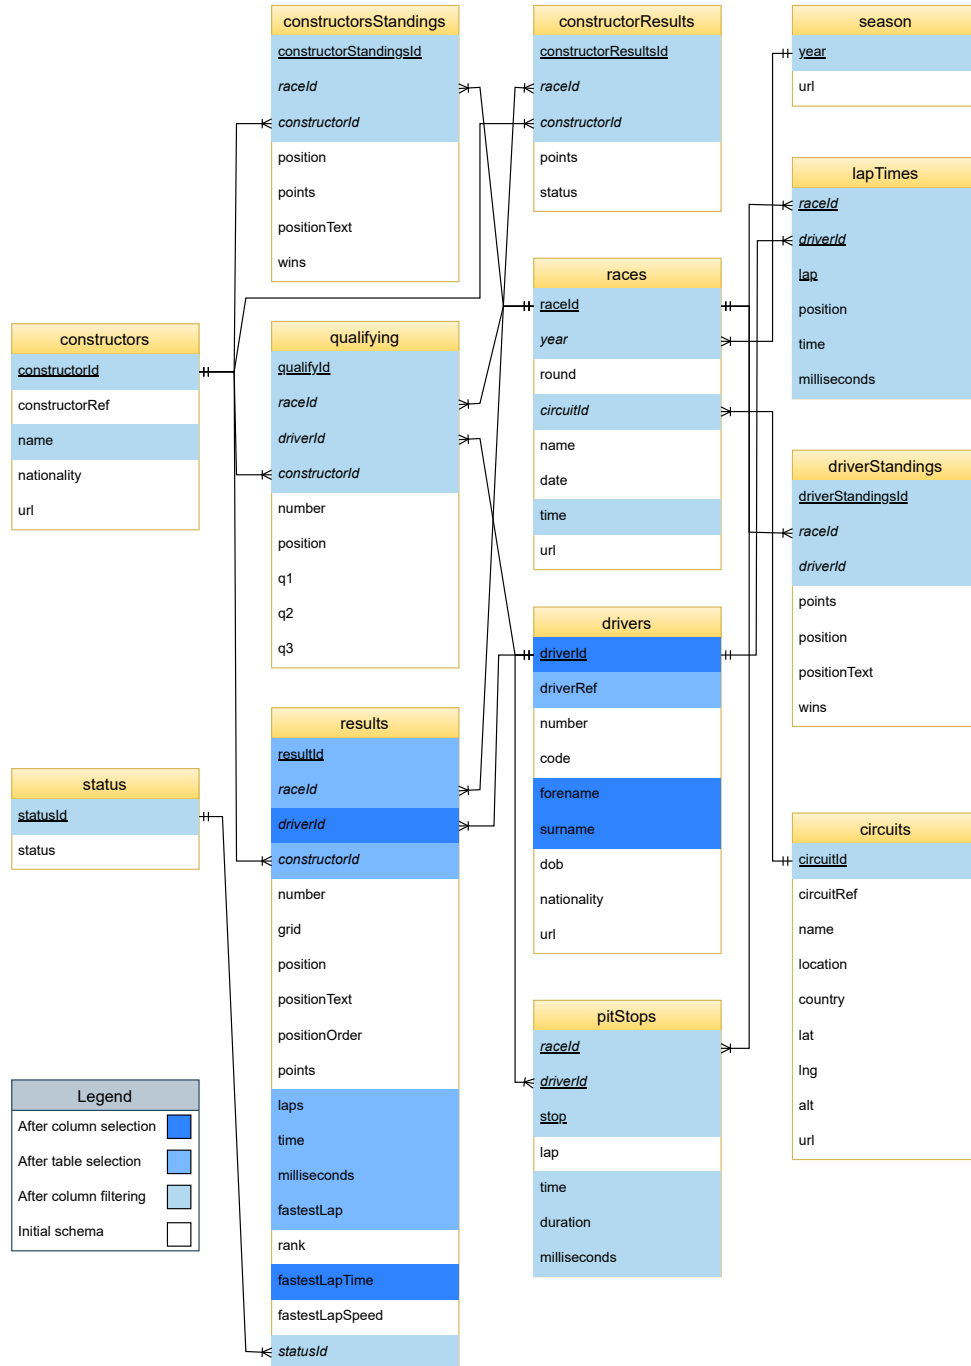




Figure 13: Schema selection example formula_1_926

G Error Analysis

To analyze our failure cases, we subsampled 147 questions from the development set of BIRD (SDS) and processed these questions using our pipeline and a vanilla GPT-4 baseline. The vanilla GPT-4 baseline replicates the GPT-4 approach from BIRD, where the question, evidence, and the full schema with all tables and columns are provided to GPT-4 with chain-of-thought reasoning prompts. In this context, the evidence refers to the hint provided alongside some questions in the dataset.



(a) Vanilla GPT4



(b) CHES

Figure 14: Distribution of Errors on Sampled Dev Set.

Figure 14 shows the categories of errors and their percentages for our approach and the baseline. "Incorrectly Predicted SQL" refers to failures in our pipeline that lead to incorrect final SQL, while ambiguous questions and incorrect golden SQL indicate problems with the dataset, where the data formats expected from the questions were ambiguous, and the golden SQL provided is incorrect.

For the vanilla baseline, 57.1% of the incorrectly predicted SQL results from incorrect schema linking issues, contributing 26.0% of the errors due to wrong columns in SELECT or JOIN. In comparison, our approach has 42.9% incorrect SQL, with only 5.4% of the errors attributed to incorrect column linking. These figures show that our approach distributes error types more evenly, indicating improvements in all categories of potential errors.

G.1 Incorrectly Predicted SQL

G.1.1 Column:

Description: The columns used in the predicted SQL are either missing or incorrect.

Question	Which drivers born after 1975 have been ranked 2? Please give their forenames and surnames.
Evidence	born after 1975 refers to <code>year(dob) > 1975</code> ;
Gold SQL	<pre>SELECT T2.forename, T2.surname FROM results AS T1 INNER JOIN drivers AS T2 ON T1.driverId = T2.driverId WHERE STRFTIME('%Y', T2.dob) > '1975' AND T1.rank = 2</pre>
Predicted SQL	<pre>SELECT T1.forename, T1.surname FROM drivers AS T1 INNER JOIN driverStandings AS T2 ON T1.driverId = T2.driverId WHERE T2.position = 2 AND STRFTIME('%Y', T1.dob) > '1975'</pre>
Error Explanation	"rank" is the correct column.
Question	Find all cards illustrated by Stephen Daniel and describe the text of the ruling of these cards. State if these cards have missing or degraded properties and values.
Evidence	cards have missing or degraded properties and value refers to <code>hasContentWarning = 1</code> ; 'Stephen Daniele' is artist;
Gold SQL	<pre>SELECT T1.id, T2.text, T1.hasContentWarning FROM cards AS T1 INNER JOIN rulings AS T2 ON T1.uuid = T2.uuid WHERE T1.artist = 'Stephen Daniele'</pre>
Predicted SQL	<pre>SELECT T1.uuid, T2.text, T1.hasContentWarning FROM cards AS T1 INNER JOIN rulings AS T2 ON T1.uuid = T2.uuid WHERE T1.artist = 'Stephen Daniele'</pre>
Error Explanation	"id" is the correct column.

G.1.2 Table:

Description: The tables used in the predicted SQL are either missing or incorrect.

Question	Among the schools with the average score in Math over 560 in the SAT test, how many schools are directly charter-funded?
Gold SQL	<pre>SELECT COUNT(T2.School Code) FROM satscores AS T1 INNER JOIN frpm AS T2 ON T1.cds = T2.CDSCode WHERE T1.AvgScrMath > 560 AND T2.Charter Funding Type = 'Directly funded'</pre>
Predicted SQL	<pre>SELECT COUNT(T1.cds) FROM satscores AS T1 INNER JOIN schools AS T2 ON T1.cds = T2.CDSCode WHERE T1.AvgScrMath > 560 AND T2.FundingType = 'Charter'</pre>
Error Explanation	"frpm" is the correct table with the correct column.
Question	Who placed the order with the id 32423?
Gold SQL	<pre>SELECT T3.client_id FROM order AS T1 INNER JOIN account AS T2 ON T1.account_id = T2.account_id INNER JOIN client AS T3 ON T2.district_id = T3.district_id WHERE T1.order_id = 32423</pre>
Predicted SQL	<pre>SELECT T2.client_id FROM order AS T1 INNER JOIN disp AS T2 ON T1.account_id = T2.account_id WHERE T1.order_id = 32423</pre>
Error Explanation	"account" and "client" are the correct tables with the correct columns.

G.1.3 Entity:

Description: The entity used for filtering is incorrect in the predicted SQL.

Question	For the customers who paid in the euro, what is their average total price of the transactions?
Gold SQL	<pre>SELECT AVG(T1.Price) FROM transactions_1k AS T1 INNER JOIN gasstations AS T2 ON T1.GasStationID = T2.GasStationID INNER JOIN customers AS T3 ON T1.CustomerID = T3.CustomerID WHERE T3.Currency = 'EUR'</pre>
Predicted SQL	<pre>SELECT AVG(T1.Price) FROM transactions_1k T1 JOIN customers T2 ON T1.CustomerID = T2.CustomerID WHERE T2.Currency = 'Euro'</pre>

Error Explanation	"EUR" is the correct entity for filtering.
Question	List out the clients who choose statement of issuance after transaction are Dispo- nent?
Gold SQL	<pre>SELECT T2.client_id FROM account AS T1 INNER JOIN disp AS T2 ON T1.account_id = T2.account_id WHERE T1.frequency = 'POPLATEK PO OBRATU' AND T2.type = 'DISPONENT'</pre>
Predicted SQL	<pre>SELECT T2.client_id FROM account AS T1 INNER JOIN disp AS T2 ON T1.account_id = T2.account_id WHERE T1.frequency = 'POPLATEK PO OBRATU' AND T2.type = 'Disponent'</pre>
Error Explanation	"DISPONENT" is the correct entity for filtering.

G.1.4 Revision:

Description: The revision node incorrectly changes the correct predicted SQL into an incorrectly predicted SQL.

Question	In Los Angeles how many schools have more than 500 free meals but less than 700 free or reduced price meals for K-12?
Gold SQL	<pre>SELECT COUNT(CDSCode) FROM frpm WHERE County Name = 'Los Angeles' AND Free Meal Count (K-12) > 500 AND 'Free Meal Count (K-12)' < 700</pre>
Predicted SQL	<pre>SELECT COUNT(CDSCode) FROM frpm WHERE Free Meal Count (K-12) > 500 AND 'FRPM Count (K-12)' < 700 AND County Name = 'Los Angeles'</pre>
Error Explanation	"Free Meal Count (K-12)" is the correct column for filtering.
Question	Find and list the names of sets which doesn't have Japanese translation but have Korean translation.
Gold SQL	<pre>SELECT name FROM sets WHERE code IN (SELECT setCode FROM set_translations WHERE language = 'Korean' AND language NOT LIKE 'Japanese')</pre>

Predicted SQL	<pre> SELECT T1.name FROM sets AS T1 WHERE EXISTS (SELECT 1 FROM set_translations AS T2 WHERE T1.code = T2.setCode AND T2.language = 'Korean') AND NOT EXISTS (SELECT 1 FROM set_translations AS T3 WHERE T1.code = T3.setCode AND T3.language LIKE 'Japanese') </pre>
Error Explanation	Incorrect logic in subquery.

G.1.5 Malformed Date:

Description: The date format in the predicted SQL is incorrect.

Question	Which year recorded the most consumption of gas paid in CZK?
Evidence	The first 4 strings of the values in the table yearmonth can represent year.
Gold SQL	<pre> SELECT SUBSTRING(T2.Date, 1, 4) FROM customers AS T1 INNER JOIN yearmonth AS T2 ON T1.CustomerID = T2.CustomerID WHERE T1.Currency = 'CZK' GROUP BY SUBSTRING(T2.Date, 1, 4) ORDER BY SUM(T2.Consumption) DESC LIMIT 1 </pre>
Predicted SQL	<pre> SELECT strftime('%Y', T2.Date) AS Year FROM yearmonth T2 JOIN customers T1 ON T2.CustomerID = T1.CustomerID WHERE T1.Currency = 'CZK' GROUP BY Year ORDER BY SUM(T2.Consumption) DESC LIMIT 1 </pre>
Error Explanation	Incorrect year formatting.

G.1.6 Evidence:

Description: The evidence was not used correctly in the predicted SQL.

Question	Name schools in Riverside which the average of average math score for SAT is greater than 400, what is the funding type of these schools?
Evidence	Average of average math = sum(average math scores) / count(schools).

Gold SQL	<pre>SELECT T1.sname, T2.Charter Funding Type` FROM satscores AS T1 INNER JOIN frpm AS T2 ON T1.cds = T2.CDSCode WHERE T2.District Name LIKE 'Riverside%' GROUP BY T1.sname, T2.Charter Funding Type` HAVING CAST(SUM(T1.AvgScrMath) AS REAL) / COUNT(T1.cds) > 400</pre>
Predicted SQL	<pre>SELECT DISTINCT T1.CDSCode, T1.FundingType FROM schools AS T1 INNER JOIN satscores AS T2 ON T1.CDSCode = T2.cds WHERE T1.County = 'Riverside' AND T2.AvgScrMath > 400</pre>
Error Explanation	Did not use the equation in the evidence.

Question	What is the detailed content of the comment of the post which was created on 7/19/2010 7:37:33 PM?
-----------------	--

Evidence	detailed content of the comment refers to Text; created on 7/19/2010 7:37:33 PM CreationDate = 2010-07-19 19:37:33.0
-----------------	---

Gold SQL	<pre>SELECT T1.Text FROM comments AS T1 INNER JOIN posts AS T2 ON T1.PostId = T2.Id WHERE T1.CreationDate = '2010-07-19 19:37:33.0'</pre>
-----------------	---

Predicted SQL	<pre>SELECT T2.Text FROM posts AS T1 INNER JOIN comments AS T2 ON T1.Id = T2.PostId WHERE T1.CreationDate = '2010-07-19 19:37:33'</pre>
----------------------	---

Error Explanation	Did not use the time formatting in the evidence.
--------------------------	--

G.1.7 NULL/DISTINCT/GROUP BY:

Description: Missing NULL, DISTINCT, or GROUP BY in the predicted SQL.

Question	When was the oldest mythic card released and what are its legal play formats?
Evidence	the oldest card refers to MIN(originalReleaseDate); mythic card refers to rarity = 'mythic'; legal play refers to status = 'legal'; play format refers to format
Gold SQL	<pre>SELECT T1.originalReleaseDate, T2.format FROM cards AS T1 INNER JOIN legalities AS T2 ON T1.uuid = T2.uuid WHERE T1.rarity = 'mythic' AND T1.originalReleaseDate IS NOT NULL AND T2.status = 'Legal' ORDER BY T1.originalReleaseDate LIMIT 1</pre>

Predicted SQL	<pre>SELECT T1.originalReleaseDate, T2.format FROM cards AS T1 INNER JOIN legalities AS T2 ON T1.uuid = T2.uuid WHERE T1.rarity = 'mythic' AND T2.status = 'Legal' ORDER BY T1.originalReleaseDate ASC LIMIT 1</pre>
----------------------	--

Error Explanation	Did not check if the selected column has NULL values.
--------------------------	---

Question	Name the foreign name of the card that has abzan watermark? List out the type of this card.
-----------------	---

Evidence	
-----------------	--

Gold SQL	<pre>SELECT DISTINCT T1.name, T1.type FROM cards AS T1 INNER JOIN foreign_data AS T2 ON T2.uuid = T1.uuid WHERE T1.watermark = 'abzan'</pre>
-----------------	--

Predicted SQL	<pre>SELECT T2.name, T2.type FROM cards AS T1 INNER JOIN foreign_data AS T2 ON T1.uuid = T2.uuid WHERE T1.watermark = 'abzan'</pre>
----------------------	---

Error Explanation	Did not use the DISTINCT keyword for the selected columns.
--------------------------	--

G.1.8 Operation:

Description: Incorrect comparison or division operations in predicted SQL.

Question	Among the male patients who have a normal level of white blood cells, how many of them have an abnormal fibrinogen level?
-----------------	---

Evidence	male patients refers to Sex = 'M'; normal level of white blood cells refers to WBC between 3.5 and 9.0; abnormal fibrinogen level refers to FG <= 150 or FG >= 450; Don't compute repetitive ones.
-----------------	--

Gold SQL	<pre>SELECT COUNT(DISTINCT T1.ID) FROM Patient AS T1 INNER JOIN Laboratory AS T2 ON T1.ID = T2.ID WHERE T2.FG <= 150 OR T2.FG >= 450 AND T2.WBC > 3.5 AND T2.WBC < 9.0 AND T1.SEX = 'M'</pre>
-----------------	---

Predicted SQL	<pre>SELECT COUNT(DISTINCT T1.ID) FROM Patient AS T1 INNER JOIN Laboratory AS T2 ON T1.ID = T2.ID WHERE T1.SEX = 'M' AND T2.WBC BETWEEN 3.5 AND 9.0 AND (T2.FG < 150 OR T2.FG > 450)</pre>
----------------------	--

Error Explanation	Incorrect comparison operation.
--------------------------	---------------------------------

Question	Among all chemical compounds identified in the database, what percent of compounds form a triple-bond.
Evidence	triple bond refers to bond_type = '#';
Gold SQL	<pre>SELECT CAST(COUNT(CASE WHEN T.bond_type = '#' THEN T.bond_id ELSE NULL END) AS REAL) * 100 / COUNT(T.bond_id) FROM bond AS T</pre>
Predicted SQL	<pre>SELECT CAST(COUNT(DISTINCT T1.molecule_id) AS REAL) * 100 / (SELECT COUNT() FROM molecule) AS percentage FROM bond AS T1 WHERE T1.bond_type = '#'</pre>
Error Explanation	Incorrect division operation with the wrong column.

G.1.9 SELECT:

Description: Incorrect formatting in SELECT in the predicted SQL.

Question	Provide the full name and email address of the Student_Club's Secretary.
Evidence	full name refers to first_name, last_name; 'Secretary' is a position of Student Club
Gold SQL	<pre>SELECT first_name, last_name, email FROM member WHERE position = 'Secretary'</pre>
Predicted SQL	<pre>SELECT first_name ' ' last_name AS full_name, email FROM member WHERE position = 'Secretary'</pre>
Error Explanation	Incorrect column formatting in SELECT.

G.2 Vague Question

G.2.1 Column:

Description: The question and evidence do not specify which columns should be included in the SELECT statement, making it impossible to determine the correct columns from the provided information and database. This leads to incorrect or additional columns in the gold SQL.

Question	What are the cards for set OGW? State the colour for these cards.
Evidence	set OGW refers to setCode = 'OGW';
Gold SQL	<pre>SELECT id, colors FROM cards WHERE id IN (SELECT id FROM set_translations WHERE setCode = 'OGW')</pre>
Predicted SQL	<pre>SELECT name, colors FROM cards WHERE setCode = 'OGW'</pre>
Error Explanation	Unclear if the question is asking for name or id.

Question	Which of these players performs the best in crossing actions, Alexis, Ariel Borysiuk or Arouna Kone?
Evidence	player who perform best in crossing actions refers to MAX(crossing);
Gold SQL	<pre>SELECT t1.player_name, t2.crossing FROM Player AS t1 INNER JOIN Player_Attributes AS t2 ON t1.player_api_id = t2.player_api_id WHERE t1.player_name IN ('Alexis', 'Ariel Borysiuk', 'Arouna Kone') ORDER BY t2.crossing DESC LIMIT 1</pre>
Predicted SQL	<pre>SELECT T1.player_name FROM Player AS T1 INNER JOIN Player_Attributes AS T2 ON T1.player_fifa_api_id = T2.player_fifa_api_id WHERE T1.player_name IN ('Alexis', 'Ariel Borysiuk', 'Arouna Kone') ORDER BY T2.crossing DESC LIMIT 1</pre>
Error Explanation	Unclear from the question if the output should include both the names and the crossing actions.

G.3 Incorrect Golden SQL

G.3.1 Evidence:

Description: The gold SQL does not use the evidence correctly.

Question	What is the percentage of Story Spotlight cards that also have a text box? List them by their ID.
Evidence	Story Spotlight cards that do not have a text box refers to isStorylight = 1 and isTextless = 1; Percentage refer to DIVIDE(SUM(count(id) where isStorylight = 1), SUM(count(id))) * 100
Gold SQL	<pre>SELECT CAST(SUM(CASE WHEN isTextless = 0 THEN 1 ELSE 0 END) AS REAL) 100 / COUNT(id) FROM cards WHERE isStorySpotlight = 1</pre>
Predicted SQL	<pre>SELECT CAST(SUM(CASE WHEN isStorySpotlight = 1 AND isTextless = 0 THEN 1 ELSE 0 END) AS REAL) 100 / COUNT(id) AS percentage FROM cards</pre>
Error Explanation	The filtering should be done in the numerator as shown in the evidence.
Question	Who is the illustrator that illustrated the least amount of cards? List the format of play of the cards that he/she illustrated.
Evidence	format of the cards refers to format; illustrator refers to artist; the least amount of cards refers to MIN(artist)

Gold SQL	<pre> SELECT T1.artist, T2.format FROM cards AS T1 INNER JOIN legalities AS T2 ON T2.uuid = T1.uuid GROUP BY T1.artist ORDER BY COUNT(T1.id) ASC LIMIT 1 </pre>
Predicted SQL	<pre> SELECT T1.artist, T2.format FROM cards AS T1 INNER JOIN legalities AS T2 ON T1.uuid = T2.uuid GROUP BY T1.artist, T2.format ORDER BY COUNT(T1.artist) ASC LIMIT 1 </pre>
Error Explanation	The gold SQL did not follow the evidence instruction to use the right column in MIN(artist).

G.3.2 Description:

Description: The gold SQL does not utilize the column or value descriptions provided in the database, leading to incorrect SQL.

Question	Which gas station has the highest amount of revenue?
Gold SQL	<pre> SELECT GasStationID FROM transactions_1k GROUP BY GasStationID ORDER BY SUM(Price) DESC LIMIT 1 </pre>
Predicted SQL	<pre> SELECT T2.GasStationID FROM gasstations AS T1 INNER JOIN transactions_1k AS T2 ON T1.GasStationID = T2.GasStationID WHERE T2.Amount IS NOT NULL AND T2.Price IS NOT NULL GROUP BY T2.GasStationID ORDER BY SUM(T2.Amount * T2.Price) DESC LIMIT 1 </pre>
Error Explanation	The description in the database indicates revenue = amount*price.
Question	What's the finish time for the driver who ranked second in 2008's Australian Grand Prix?
Gold SQL	<pre> SELECT T1.time FROM results AS T1 INNER JOIN races AS T2 ON T1.raceId = T2.raceId WHERE T1.rank = 2 AND T2.name = 'Australian Grand Prix' AND T2.year = 2008 </pre>

Predicted SQL	<pre> SELECT T2.time FROM races AS T1 INNER JOIN results AS T2 ON T1.raceId = T2.raceId WHERE T1.name = 'Australian Grand Prix' AND T1.year = 2008 AND T2.positionOrder = 2 </pre>
Error Explanation	The description in the database indicates the position order is the finishing order of positions while rank is the starting rank positioned by fastest lap speed.

G.3.3 Aggregation:

Description: The gold SQL misses aggregation operations.

Question	What is the amount of debt that client number 992 has, and how is this client doing with payments?
Gold SQL	<pre> SELECT T3.amount, T3.status FROM client AS T1 INNER JOIN account AS T2 ON T1.district_id = T2.district_id INNER JOIN loan AS T3 ON T2.account_id = T3.account_id WHERE T1.client_id = 992 </pre>
Predicted SQL	<pre> SELECT SUM(T3.amount) AS total_debt, T3.status FROM client AS T1 INNER JOIN disp AS T2 ON T1.client_id = T2.client_id INNER JOIN loan AS T3 ON T2.account_id = T3.account_id WHERE T1.client_id = 992 GROUP BY T3.status </pre>
Error Explanation	Gold SQL did not use SUM to sum the full amount.
Question	How many cards have frame effect as extendedart? List out the id of those cards.
Gold SQL	<pre> SELECT id FROM cards WHERE frameEffects = 'extendedart' GROUP BY id </pre>
Predicted SQL	<pre> SELECT COUNT(id) FROM cards WHERE frameEffects = 'extendedart' UNION ALL SELECT id FROM cards WHERE frameEffects = 'extendedart' </pre>
Error Explanation	The gold SQL did not return a count as asked by the question.

G.3.4 Column:

Description: The gold SQL does not have all the expected columns from the question in SELECT.

Question	List out all post that are related to post ID 61217 and what is the popularity of this post?
Evidence	post related refers to RelatedPostId; popularity refers to ViewCount
Gold SQL	<pre>SELECT T1.ViewCount FROM posts AS T1 INNER JOIN postLinks AS T2 ON T1.Id = T2.PostId WHERE T2.PostId = 61217</pre>
Predicted SQL	<pre>SELECT T2.RelatedPostId, T1.ViewCount FROM posts AS T1 INNER JOIN postLinks AS T2 ON T1.Id = T2.PostId WHERE T2.PostId = 61217</pre>
Error Explanation	The gold SQL did not select all the columns asked by the question.
Question	List down the bond type for molecules from molecule id TR000 to TR050.
Evidence	double bond refers to bond_type = '='; single bond refers to bond_type = '-'; triple bond refers to bond_type = '#';
Gold SQL	<pre>SELECT T2.molecule_id, T2.bond_type FROM molecule AS T1 INNER JOIN bond AS T2 ON T1.molecule_id = T2.molecule_id WHERE T1.molecule_id BETWEEN 'TR000' AND 'TR050'</pre>
Predicted SQL	<pre>SELECT T2.bond_type FROM molecule AS T1 INNER JOIN bond AS T2 ON T1.molecule_id = T2.molecule_id WHERE T1.molecule_id LIKE 'TR0%' AND T1.molecule_id BETWEEN 'TR000' AND 'TR050'</pre>
Error Explanation	The gold SQL has more columns selected than what the question asked for.

G.3.5 Filtering:

Description: The gold SQL does not filter using the correct entity or is missing necessary filters.

Question	How many male customers who are living in North Bohemia have average salary greater than 8000?
Evidence	Male means that gender = 'M'; A3 refers to region; A11 pertains to average salary.
Gold SQL	<pre>SELECT COUNT(T1.client_id) FROM client AS T1 INNER JOIN district AS T2 ON T1.district_id = T2.district_id WHERE T1.gender = 'M' AND T2.A3 = 'North Bohemia' AND T2.A11 > 8000</pre>

Predicted SQL	<pre> SELECT COUNT(T1.client_id) FROM client AS T1 INNER JOIN district AS T2 ON T1.district_id = T2.district_id WHERE T1.gender = 'M' AND T2.A3 = 'north Bohemia' AND T2.A11 > 8000 </pre>
Error Explanation	The correct entity name has a lower case in north.
Question	How many users from New York have a teacher and supporter badge?
Evidence	"Supporter" and "Teachers" are both Name of badge; 'New York' is the Location; user refers to UserId
Gold SQL	<pre> SELECT COUNT(DISTINCT T1.Id) FROM badges AS T1 INNER JOIN users AS T2 ON T1.UserId = T2.Id WHERE T1.Name IN ('Supporter', 'Teacher') AND T2.Location = 'New York' </pre>
Predicted SQL	<pre> SELECT COUNT(DISTINCT T1.Id) FROM users AS T1 INNER JOIN badges AS T2 ON T1.Id = T2.UserId WHERE T1.Location = 'New York' AND T2.Name IN ('Supporter', 'Teacher') GROUP BY T1.Id HAVING COUNT(DISTINCT T2.Name) = 2 </pre>
Error Explanation	The gold SQL missed a filtering that indicates a user is both teacher and a supporter.

G.3.6 NULL/DISTINCT/GROUP BY:

Description: The gold SQL missed checking NULL, adding DISTINCT, or adding GROUP BY.

Question	Calculate the total amount spent on speaker gifts and list the name of the event they were spent on.
Evidence	total amount spent = SUM(spent) where category = 'Speaker Gifts'
Gold SQL	<pre> SELECT SUM(T1.spent), T2.event_name FROM budget AS T1 INNER JOIN event AS T2 ON T1.link_to_event = T2.event_id WHERE T1.category = 'Speaker Gifts' </pre>
Predicted SQL	<pre> SELECT SUM(T2.spent), T1.event_name FROM event AS T1 INNER JOIN budget AS T2 ON T1.event_id = T2.link_to_event WHERE T2.category = 'Speaker Gifts' GROUP BY T1.event_name </pre>
Error Explanation	The gold SQL misses GROUP BY.