



Pervasive Annotation Errors Break Text-to-SQL Benchmarks and Leaderboards

Tengjun Jin
University of Illinois (UIUC)
Urbana, USA
tengjun2@illinois.edu

Yuxuan Zhu
University of Illinois (UIUC)
Urbana, USA
yxx404@illinois.edu

Yoojin Choi
University of Illinois (UIUC)
Urbana, USA
yoojinc3@illinois.edu

Daniel Kang
University of Illinois (UIUC)
Urbana, USA
ddkang@illinois.edu

ABSTRACT

Researchers have proposed numerous text-to-SQL techniques to streamline data analytics and accelerate the development of data-driven applications. To compare these techniques and select the best one for deployment, the community depends on public benchmarks and their leaderboards. Since these benchmarks heavily rely on human annotations during question construction and answer evaluation, the validity of the annotations is crucial.

In this paper, we conduct an empirical study that (i) benchmarks annotation error rates for two widely used text-to-SQL benchmarks, BIRD and Spider 2.0-Snow, and (ii) corrects a subset of the BIRD development (Dev) set to measure the impact of annotation errors on text-to-SQL agent performance and leaderboard rankings. Through expert analysis, we show that BIRD Mini-Dev and Spider 2.0-Snow have error rates of 52.8% and 62.8%, respectively. We re-evaluate all 16 open-source agents from the BIRD leaderboard on both the original and the corrected BIRD Dev subsets. We show that performance changes range from -7% to 31% (in relative terms) and rank changes range from -9 to +9 positions. We further assess whether these impacts generalize to the full BIRD Dev set. We find that the rankings of agents on the uncorrected subset correlate strongly with those on the full Dev set (Spearman’s $r_s=0.85$, $p=3.26e-5$), whereas they correlate weakly with those on the corrected subset (Spearman’s $r_s=0.32$, $p=0.23$). These findings show that annotation errors can significantly distort reported performance and rankings, potentially misleading research directions or deployment choices.

PVLDB Reference Format:

Tengjun Jin, Yoojin Choi, Yuxuan Zhu, and Daniel Kang. Pervasive Annotation Errors Break Text-to-SQL Benchmarks and Leaderboards. PVLDB, 19(5): 931 - 944, 2026. doi:10.14778/3796195.3796206

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at https://github.com/uiuc-kang-lab/text_to_sql_benchmarks.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Proceedings of the VLDB Endowment, Vol. 19, No. 5 ISSN 2150-8097. doi:10.14778/3796195.3796206

1 INTRODUCTION

As text-to-SQL techniques become increasingly important for data analytics and data-driven applications [20, 21, 43, 44, 47, 54, 60], researchers and practitioners have introduced a variety of text-to-SQL benchmarks [13, 24, 29, 30, 34, 35, 48, 61, 63, 65]. These benchmarks guide researchers in improving text-to-SQL techniques and help practitioners select the best text-to-SQL technique for downstream applications in real-world scenarios.

Unfortunately, current text-to-SQL benchmarks are unreliable due to annotation errors. As reported in prior work, BIRD [34] has annotation errors in 32% of examples in the mini development set (Mini-Dev) [4] and 49% of financial domain examples [59]. Such errors can distort reported performance and rankings of agents and mislead researchers and practitioners.

To understand how we can build reliable text-to-SQL benchmarks with minimal annotation errors, we conduct an in-depth analysis of annotation errors in text-to-SQL benchmarks. We design our analysis based on four recurring text-to-SQL annotation error patterns. Given a target database \mathcal{D} , human annotators curate text-to-SQL examples, each comprising a natural language input \mathcal{T} and a ground-truth SQL query Q .¹ Annotation errors can arise within any individual component or from inconsistencies across components, including:

- E1. Mismatches between the semantics of Q and the logic of \mathcal{T} .
- E2. Mismatches between the semantics of Q and \mathcal{D} due to a limited understanding of the data or the schema.
- E3. Mismatches between the semantics of Q and the domain knowledge relevant to \mathcal{T} , or misannotated domain knowledge in \mathcal{T} .
- E4. Ambiguity in \mathcal{T} , such as multiple possible interpretations or an unclear output format.

Figure 1a illustrates a misuse of a Snowflake function (E1). Figure 1b presents an incorrect annotation from BIRD due to the annotator’s lack of domain knowledge (E3). We present additional representative examples for each error pattern in Section 4.4.

We conducted a human-in-the-loop, three-stage audit, with all final decisions made by human SQL experts, to benchmark annotation error rates for two widely used benchmarks, BIRD [34] and Spider 2.0-Snow [30]. To facilitate the audit, we developed SAR-Agent (SQL Annotation Reviewer agent), the first AI agent that assists SQL experts in detecting annotation errors in text-to-SQL

¹Input \mathcal{T} consists of a user question and, optionally, external knowledge.

Annotated user question: Can you provide a daily weather summary for July 2019 within a 5 km radius of latitude 26.75 and longitude 51.5? ...

Annotated SQL

```
... ST_DWITHIN(ST_GEOGFROMWKB("TRI"."geography"),  
ST_POINT(26.75, 51.5), 5000) ...
```

Issue: Snowflake’s ST_POINT expects (longitude, latitude). The query passes ST_POINT(26.75, 51.5), which inverts the order.

(a) sf_bq291 of Spider 2.0-Snow. Annotators mistakenly swapped the longitude–latitude order in the ST_POINT arguments (E1).

Annotated user question: Which state special schools have the highest number of enrollees from grades 1 through 12?

Annotated external knowledge: State Special Schools refers to DOC = 31; Grades 1 through 12 means K-12

Annotated SQL

```
SELECT T2.School  
FROM frpm AS T1 INNER JOIN schools AS T2  
ON T1.CDSCode = T2.CDSCode WHERE T2.DOC = 31  
ORDER BY T1."Enrollment_(K-12)" DESC LIMIT 1
```

Issue: ‘Enrollment (K-12)’ includes kindergarten and is therefore not equivalent to enrollment for grades 1 through 12.

(b) Example 46 of BIRD Mini-Dev. Annotators had limited understanding of “K-12” (E3).

Figure 1: Two examples from existing text-to-SQL benchmarks that demonstrate incorrect annotations.

benchmarks. Our three-stage audit consisted of: (1) SAR-Agent producing per-annotation diagnostic reports; (2) human SQL experts adjudicating all agent-flagged cases, verifying the stated error reasons and assigning final correctness labels; and (3) human SQL experts extending manual review to additional examples based on the recurring error patterns.

We manually corrected a subset of the BIRD Dev set to measure the impact of annotation errors on text-to-SQL agent performance and leaderboard rankings. To support manual corrections by SQL experts, we introduced SAPAR (SQL Annotation Pipeline with an AI Agent Reviewer), which integrated SAR-Agent into the standard text-to-SQL annotation pipeline [34]. We randomly sampled 100 of the 1,534 BIRD Dev examples and, guided by SAPAR, manually corrected all identified errors. We then re-evaluated all 16 open-source text-to-SQL agents from the BIRD leaderboard on both the original and the corrected Dev subsets.

Below, we (i) quantify annotation error rates in BIRD Mini-Dev (an official subset of the BIRD Dev set) [34] and Spider 2.0-Snow [30], (ii) analyze their impact on agent performance and leaderboard rankings, and (iii) evaluate SAR-Agent’s error-detection capability to demonstrate its utility for improving benchmark quality.

High error rates in existing text-to-SQL benchmarks. Our analysis reveals a wider range of annotation errors in text-to-SQL benchmarks than previously recognized. We executed SAR-Agent on BIRD Mini-Dev and Spider 2.0-Snow to generate diagnostic reports for each example. After manual verification, we find that 52.8% of the examples in BIRD Mini-Dev [34] contain annotation

errors, which is 20.5% higher than the previously reported rate of 32.3% [4]. In addition, we are the first to comprehensively examine annotation errors in Spider 2.0-Snow [30], a recently released benchmark and successor to the widely used Spider 1.0 [65]. The Spider 2.0 team applied large language models (LLMs) to paraphrase user questions for clarity and corrected 45% of examples in the first validation round and 5% in the second [30]. Nevertheless, across all 121 problems for which ground-truth SQL is publicly available, we still identify an annotation error rate of 62.8%.

The unreliability of the BIRD leaderboard. These annotation errors cause severe misestimation and misranking of agents’ performance. Through our re-evaluation of 16 agents on both the original and the corrected BIRD Dev subsets, we identify significant performance changes ranging from -7% to 31% in relative terms and ranking position changes ranging from -9 to $+9$. For instance, we find the performance of CHESS [54], an agent previously ranked 7th among our 16 selected agents, increases from 62% to 81% after corrections and moves to 1st place.

We further computed two Spearman’s rank correlation coefficients (r_s) to assess the impact of annotation errors on agent rankings. We find that the rankings on the original Dev subset correlate strongly with those on the original full Dev set ($r_s=0.85$, $p=3.26e-5$), but have a weak correlation with the rankings on the corrected Dev subset ($r_s=0.32$, $p=0.23$).

These results demonstrate that annotation errors can distort reported performance and rankings, undermining the reliability of the leaderboard. Such unreliability can potentially misguide researchers and practitioners in improving text-to-SQL techniques and picking the best agent.

The effectiveness of SAR-Agent. SAR-Agent demonstrates its effectiveness in detecting annotation errors. Our manual verification of diagnostic reports for each annotated example shows that SAR-Agent achieves 83% precision on BIRD Mini-Dev. On Spider 2.0-Snow, for which no prior analyses exist, SAR-Agent achieves 89% precision. We further compared SAR-Agent’s detected error examples (after manual verification) on BIRD Mini-Dev with those reported by the Arcwise team’s SQL experts [4]. Among the reported 161 error examples, SAR-Agent attains an 85.7% hit rate. SAR-Agent also identifies 41.6% more error examples than Arcwise.

We summarize our contribution as follows:

- (1) Through an in-depth manual review by SQL experts, we find annotation error rates of 52.8% in BIRD Mini-Dev and 62.8% in Spider 2.0-Snow.
- (2) By re-evaluating all 16 open-source text-to-SQL agents from the BIRD leaderboard, we identify substantial performance shifts, with relative performance changes ranging from -7% to 31% and ranking changes ranging from -9 to $+9$.
- (3) We introduce a toolkit, comprising SAR-Agent and SAPAR, for annotation error detection and correction.

2 EXPERIMENT DESIGN

In this work, we develop a toolkit to help human SQL experts analyze existing text-to-SQL benchmarks and generate insights from them. We design experiments to investigate the following research questions:




Figure 2: Number of evaluated methods in the leaderboard and ground-truth SQL tokens across text-to-SQL benchmarks (data cutoff: October 20, 2025).

- Q1. What are the annotation error rates in widely used text-to-SQL benchmarks?
- Q2. How do annotation errors affect text-to-SQL agents’ performance and rankings on the leaderboard?
- Q3. How effective is SAR-Agent at detecting annotation errors?

2.1 Benchmark Selection

We selected two text-to-SQL benchmarks, BIRD [34] and Spider 2.0-Snow [30], for our empirical study based on usage and query complexity.² We quantified usage by the number of evaluated methods on each benchmark’s leaderboard, and complexity by the average number of tokens in the ground-truth SQL queries (using whitespace tokenization [30]) across six benchmarks [30, 34, 55, 56, 65]. As shown in Figure 2, BIRD has the largest number (100) of evaluated methods in its leaderboard. Spider 2.0-Snow has the most complex SQL queries (highest average token count of its ground-truth SQL queries, 161.8 tokens per SQL query) and is the successor to the widely used Spider 1.0 [65].

2.2 Toolkit for Annotation Error Detection and Correction

We developed a toolkit comprising SAR-Agent and SAPAR to facilitate manual analysis and correction of text-to-SQL annotation errors. For error detection, we ran SAR-Agent on the annotations to generate per-example diagnostic reports. SQL experts then reviewed the reported issues and assigned final correctness labels. For annotation error correction, we manually fixed the identified errors in a sampled BIRD Dev subset based on SAPAR. We provide implementation details of SAR-Agent and SAPAR in Section 3.

²BIRD note: BIRD consists of Train, Dev and Test sets. We focus on the Dev set because (i) the Train set is only used in studies that fine-tune LLMs, (ii) the Test set is hidden, and (iii) researchers rely on the Dev set for local evaluation and error analysis. Spider 2.0-Snow note: The Spider 2.0 gold queries were updated on Oct. 29, 2025. Since our evaluation was conducted before that, we used the annotations as of Aug. 20, 2025, available at: https://github.com/xlang-ai/Spider2/tree/84911f375c87903c73ed5fe93769b593a14efb42/spider2-snow/evaluation_suite/gold/sql

2.3 Evaluation Overview

We structure our experiments around our research questions. In Section 4, we use a three-step examination procedure to detect annotation errors in two widely used benchmarks, BIRD Mini-Dev [34] and Spider 2.0-Snow [30], and quantify their error rates to address Q1. In Section 5, we sample 100 examples from the BIRD Dev set and manually correct annotation errors by applying SAPAR. To answer Q2, we then re-evaluate all 16 open-source agents listed on the BIRD leaderboard on both the original and the corrected versions of this sampled subset. We quantify changes in execution accuracy and the relative ranking among these agents. In Section 6, we address Q3 by evaluating SAR-Agent’s performance and comparing the number of errors it detects with those identified in a prior audit by SQL experts from Arcwise [4].

3 TOOLKIT

In this section, we describe the implementation of our toolkit for detecting and correcting annotation errors. We first present the design of SAR-Agent, then detail the SAPAR architecture.

3.1 SAR-Agent

To support effective detection of annotation errors and assess error rates across different text-to-SQL benchmarks, we develop SAR-Agent, a SQL Annotation Reviewer agent. SAR-Agent evaluates both the correctness and potential ambiguity of annotations. In practice, data analysts typically explore databases and construct SQL queries in incremental steps, a process known as exploratory data analysis (EDA) [23, 37]. SAR-Agent mirrors this workflow by conducting multi-turn interactions with the database and verifying annotations step by step. We present an example demonstrating how SAR-Agent verifies annotations and generates diagnostic reports in the supplementary material [25].

We present the framework of SAR-Agent in Figure 3 and describe implementation details below.

Core mechanism. Rather than relying on a single-turn LLM judgment, SAR-Agent verifies each annotation incrementally. In each iteration, it formulates a targeted verification query to probe a specific aspect of the annotated SQL, executes it against the database, and inspects the result. All prior verification queries and their execution results are retained in the agent’s memory to guide the next iteration. In the final step, the agent synthesizes its findings into a comprehensive diagnostic report.

Agent inputs. SAR-Agent takes as input the following:

- (1) *Text-to-SQL annotation*: The natural language input \mathcal{T} and the corresponding ground-truth SQL query Q , both annotated by humans. \mathcal{T} comprises the user question and, when applicable, external knowledge that supplements the question.
- (2) *Schema of the underlying database \mathcal{D}* : Table and column specifications, including table names, column names, and column descriptions. Given the schema complexity of Spider 2.0-Snow (an average of 812.1 columns per database [30]), we restrict the schema to the gold tables (i.e., those referenced by the ground-truth SQL query), as provided by the Spider team.³

³The gold tables for each query can be found at <https://github.com/xlang-ai/Spider2/blob/main/methods/gold-tables/spider2-snow-gold-tables.jsonl>.




Figure 3: Overview of the SAR-Agent framework. At each iteration, the agent proposes and executes a query to validate the annotation. In the final iteration, it generates a diagnostic report that assesses the annotation’s correctness and ambiguity.

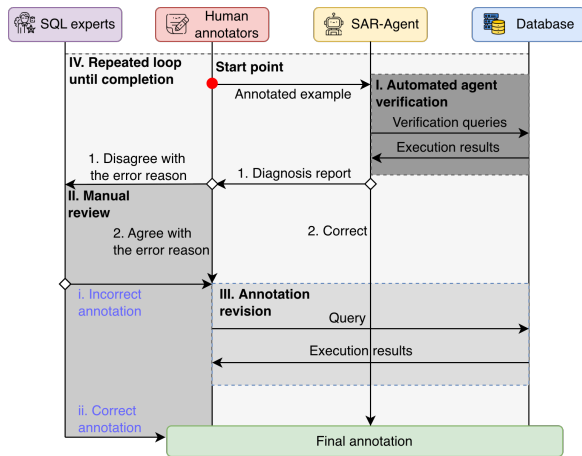


Figure 4: Overview of SAPAR. SAR-Agent audits the annotation and generates a diagnostic report for expert review. If errors are found, annotators iteratively revise the annotation until SAR-Agent or human SQL experts judge it correct.

Agent environment. The environment includes the databases for each example. To verify annotations, the agent can directly access these databases by calling predefined functions.

Agent outputs. SAR-Agent generates two files: one records the proposed verification queries, and the other presents the final diagnostic report. We describe these files as follows:

- (1) *Verification queries:* The agent logs each verification query it proposes, with accompanying explanations and the corresponding execution results.
- (2) *Final diagnostic report:* After executing verification queries, the agent generates a comprehensive diagnostic report. This report includes a judgment on question ambiguity, an assessment of SQL query correctness, detailed explanations, and a proposed corrected query if an error is detected.⁴

Agent functions. SAR-Agent is implemented using a function calling framework. We design functions that allow it to interact with the database and generate the final diagnostic report, including:

- (1) `execute_query`: executes a query on the database and returns the execution results to the agent.
- (2) `terminate`: generates a final diagnostic report and terminates.

⁴The corrected query is intended to help annotators understand the agent’s reasoning. We do not evaluate the accuracy of these corrected queries, as some examples require revisions to both the natural language question and the SQL query.

3.2 SAPAR

To enhance the efficiency and quality of annotation error correction, we propose SAPAR, which integrates SAR-Agent into the standard text-to-SQL annotation pipeline [34]. Existing pipelines require substantial human effort to ensure benchmark quality and provide rigorous evaluation. For example, BIRD [34] uses double-blind annotation followed by expert review. However, we still find pervasive annotation errors in BIRD Mini-Dev.

SAPAR augments expert review with SAR-Agent, improving both the efficiency and accuracy of annotation. First, annotators can receive annotation diagnoses within minutes, enabling immediate corrections after initial annotation rather than at the end of a batch. Second, this workflow further increases efficiency by reducing the manual review effort of SQL experts. Text-to-SQL benchmarks are typically cross-domain (e.g., BIRD contains 37 domains), making it time-consuming for reviewers to master all necessary domain knowledge. SAR-Agent can reduce human efforts by eliminating the routine need for a SQL expert to review each annotation. Third, through our evaluation results in Section 6, we find that SAR-Agent can detect more errors than the human experts in prior work [4], thereby improving annotation quality.

We show the overall architecture of SAPAR in Figure 4, and describe the workflow as follows:

Step 1: Automated agent verification. As a reviewer, SAR-Agent interacts with the database to verify the annotation and generates a diagnostic report. If the report confirms that the annotation is correct, it is accepted. Otherwise, SAR-Agent returns the report to the annotator, including any detected issues.

Step 2: Manual review and adjudication. The annotators review the reported issues. If they agree with the issues, they revise the annotation. If they disagree, they consult a SQL expert reviewer for a final decision. If the expert considers the annotation correct, it is accepted; otherwise, the annotators are required to revise it.

Step 3: Annotation revision. When SAR-Agent detects errors, human annotators revise the original annotations using the diagnostic report, which details error reasons and proposes a candidate query revision. For the SQL query, they may adopt the revision suggested by SAR-Agent after manual verification. If they disagree with the proposal, they need to revise the query themselves. For the user question and external knowledge, annotators need to manually correct any errors to ensure accuracy and natural, idiomatic phrasing. After revision, the workflow returns to Step 1.

We used SAPAR to correct annotation errors in existing text-to-SQL benchmarks. When researchers develop a new text-to-SQL

benchmark, they can apply SAPAR to improve annotation quality after creating the initial annotations.

4 BIRD AND SPIDER 2.0 HAVE PERVASIVE ANNOTATION ERRORS

We conducted a comprehensive audit of two widely studied text-to-SQL benchmarks: BIRD Mini-Dev [34] and Spider 2.0-Snow [30]. In this section, we first describe our three-step examination procedure. We then report the error rate for each benchmark. Finally, we analyze all detected annotation errors based on four patterns introduced in Section 1.

4.1 Experimental Settings

Data settings. We evaluated two widely-studied text-to-SQL benchmarks: BIRD Mini-Dev [34] and Spider 2.0-Snow [30].

- (1) *BIRD Mini-Dev*: Following prior work [4], we use the first version of BIRD Mini-Dev, which contains 500 SELECT-only text-to-SQL examples carefully chosen by the BIRD team from the original BIRD Dev set.⁵ However, prior work reports two duplicate examples, leaving 498 unique examples [4]. We find that the databases provided in the BIRD Mini-Dev set contain fewer rows compared to those in the BIRD Dev set. To better evaluate the correctness of SQL queries, we use databases from the BIRD Dev set in our experiments.
- (2) *Spider 2.0-Snow*: Spider 2.0-Snow is a variant of Spider 2.0. We choose this variant because teams at leading technology companies benchmark their agents on it rather than on Spider 2.0-lite. Spider 2.0-Snow contains a total of 547 examples, with gold queries publicly available for 121 of them. Therefore, we evaluated the correctness of these 121 examples.

Metrics. We define the error rate of each text-to-SQL benchmark as the percentage of incorrect examples in the benchmark. We consider an example to be incorrect if it exhibits any of E1 - E4.

Agent settings. We ran SAR-Agent with the OpenAI o3 model [42] and configured the agent to complete the annotation error diagnostic within a maximum of 30 iterations.⁶

4.2 Three-step Examination

Our three-step examination procedure consists of automatic detection, manual verification, and additional manual detection. We detail each step below.

Automatic detection. We first executed SAR-Agent on two text-to-SQL benchmarks: BIRD Mini-Dev [34] and Spider 2.0-Snow [30]. For each example, SAR-Agent generated a diagnostic report.

Manual verification. After generating the diagnostic report, we manually verified the correctness of error reasons for all examples that SAR-Agent labels as incorrect or ambiguous. Two authors independently assessed each example to ensure the accuracy of the verification process. In instances of disagreement, all four authors

⁵The BIRD Mini-Dev dataset is available at https://github.com/bird-bench/mini_dev/tree/main/llm/mini_dev_data.

⁶We release SAR-Agent’s prompt and scaffolding at https://github.com/uiuc-kang-lab/text_to_sql_benchmarks/tree/main/SAR-Agent.

Table 1: Error pattern distribution in BIRD Mini-Dev and Spider 2.0-Snow. Each example may contain multiple errors across different patterns.

Error pattern	# Examples	
	BIRD Mini-Dev	Spider 2.0-Snow
E1	77 (29.28%)	37 (48.68%)
E2	152 (57.79%)	44 (57.89%)
E3	28 (10.65%)	9 (11.84%)
E4	78 (29.66%)	18 (23.68%)
Examples with annotation errors	263	76

collectively reviewed and discussed the examples to reach a consensus on whether an example should be classified as incorrect. Our manual verification process consists of two main parts. First, for questions requiring domain-specific knowledge, we consulted authoritative online resources to cross-check the explanations produced by the LLM against the definitions provided by human annotators. For example, in question 46, the annotator asserted that “grades 1-12 means K-12.” However, SAR-Agent identified this interpretation as incorrect, noting that “K-12” includes kindergarten. To validate this assessment, we referred to the definition of K-12 on Wikipedia [15], which aligns with the judgment of SAR-Agent. Second, for each annotated query, we executed a set of verification queries, such as subqueries or modified queries based on the listed error reasons, on the relevant database to assess the correctness of the SAR-Agent’s judgment.

Additional manual detection. Through manual verification, we identified recurring error patterns across different examples. To further investigate annotation errors, we examined other examples within the same database to determine whether similar errors are present elsewhere [26]. For BIRD Mini-Dev, we also combined the errors identified in prior work [4].

4.3 Evaluation Results

Using the three-step examination, we identify annotation errors in 263 of the 498 (52.8%) BIRD Mini-Dev examples, including 161 examples reported in prior work [4] and additional errors uncovered by our pipeline. For Spider 2.0-Snow, we find annotation errors in 76 of 121 examples, resulting in an error rate of 62.8%.

4.4 Analysis of Annotation Errors

We begin by detailing the classification criteria for each error pattern, then present their distributions in BIRD Mini-Dev and Spider 2.0-Snow, and conclude with representative examples drawn from both benchmarks.

Classification criteria for error patterns. We detail the criteria used to classify the four error patterns introduced in Section 1. Given a natural language input \mathcal{T} , a ground-truth SQL query Q , and the corresponding database \mathcal{D} :

- E1.** We flag E1 by comparing \mathcal{T} and Q without accessing \mathcal{D} . Typical cases include missing or extraneous filter predicates in Q , incorrect comparison operators (e.g., \mathcal{T} specifies “strictly

Table 2: Representative examples for each error pattern in BIRD Mini-Dev. The question IDs in brackets indicate questions that were marked as noisy in prior work, but due to different types of errors.

Question IDs	Error Pattern	Explanation
92, 149	E1	BETWEEN ... AND ... includes the boundary values, it conflicts with the question’s requirement for exclusive comparisons.
12, 36, 40, 50, (17)	E2	The query does not include the restriction <code>rtype = ‘S’</code> .
633, 639, 640, 685, (634, 637)	E2	The queries consider users who ever edited rather than the user who owns the post.
1376, 1378, 1403	E2	No aggregation over events.
32, 46, 62	E3	“K-12” is not equivalent to grades 1 through 12.
862, 877, 881, 940, 954	E3	Drivers with a “+n Lap” status also finished the race.
1156, 1229, 1235, (1168, 1171, 1227, 1232, 1239, 1243)	E3	The age calculation is not accurate for patients.
1175	E4	“Doctor’s diagnosis” could come from either Examination.Diagnosis or Patient.Diagnosis.

greater than 100” while Q uses “ ≥ 100 ”), or an incorrect formula (e.g., \mathcal{T} asks for “% faster” but Q computes $(\text{speed853} - \text{speed854}) / \text{speed853} \times 100$).

- E2.** Detecting **E2** requires reasoning over \mathcal{D} (both schema and data). We examine the schema and data and execute queries to flag **E2**. Examples include joining on a non-unique key without the necessary deduplication, missing necessary aggregation, or ignoring required format conversions.
- E3.** We flag **E3** when the annotation relies on domain knowledge that is incorrect or inconsistent with authoritative sources. We verify such knowledge by consulting credible documentation and cross-checking references.
- E4.** We flag **E4** when the user question in \mathcal{T} admits multiple reasonable interpretations that map to different, non-equivalent SQL queries producing different results over \mathcal{D} , or when the desired output specification is underspecified.

Distribution of error patterns. Following our classification criteria, we computed the distribution of error patterns in BIRD Mini-Dev and Spider 2.0-Snow. Because an example can contain multiple errors, it may be assigned to multiple patterns. As shown in Table 1, **E2**, attributable to a limited understanding of the data or the schema, is the most common error pattern in both benchmarks. Among examples containing annotation errors, **E2** occurs in 57.79% of BIRD Mini-Dev examples and 57.89% of Spider 2.0-Snow examples.

Table 3: Representative examples for each error pattern in Spider 2.0-Snow.

Question IDs	Error Pattern	Explanation
sf_bq263, sf_bq271, sf_bq273, sf_bq294	E1	The function TO_TIMESTAMP (end_date) sets the time to the start of the end date (00:00:00.000), instead of to the end of the day (23:59:59.999) as required.
sf_bq050, sf_bq209, sf_bq455	E1	The SQL queries lack the required filters specified by the question.
sf_bq099, sf_bq248, sf_bq422, sf_local1263	E2	Using JOIN or FLATTEN operators inflates row counts; the query lacks DISTINCT for deduplication.
sf_bq052, sf_bq246	E3	Forward citations are miscalculated.
sf_bq017, sf_bq068, sf_bq099, sf_bq182, sf_bq193, sf_bq222, sf_bq223	E4	The gold SQL queries apply additional quotation mark replacements that the question does not specify.

Representative Errors in BIRD Mini-Dev. We provide representative, newly discovered examples from BIRD Mini-Dev for each of the four error patterns in Table 2 and discuss them in more detail.

Examples of E1. The annotators used the inclusive operator BETWEEN ... AND ... in the gold query, but the question specifies strict inequalities ($>$, $<$) that exclude the boundary values. Consequently, the logic of these queries is incorrect.

Examples of E2. We find three typical errors for **E2**. Firstly, in the five listed questions related to the `california_schools` database, we find the annotated queries omit the predicate `rtype = ‘S’`. Although annotators labeled the `rtype` column as “unuseful” in the table schema, our investigation on the database confirms that `rtype` is useful as it distinguishes schools from districts. Therefore, for questions referencing schools, the ground truth SQL queries should include `rtype = ‘S’`.

Secondly, in the `codebase_community` database, the `posts` table records the owner of each post, while the `postHistory` table tracks the editors. In the six listed questions, phrases such as “posted by” or “his post” should refer to the post owner, not to the editors.

Thirdly, in the `student_club` database, each event can have multiple corresponding rows in the `budget` table. Therefore, when calculating the budget for an event, the query should aggregate the budget values associated with that event.

Examples of E3. We find three typical errors for **E3**. Firstly, in the three listed examples in Table 2 from the `california_schools` database, the annotators mistakenly interpreted “K-12” as referring solely to “grades 1–12,” thereby excluding kindergarten.

Secondly, under Formula 1’s 90% rule, a driver is classified as a finisher if they complete at least 90% of the distance covered by the winner [7]. Therefore, some drivers with a “+n Lap” status, where n is no more than 10% of the total laps, should also be considered finishers. However, the annotators omitted these drivers.

Thirdly, according to the Centers for Disease Control and Prevention (CDC), a patient’s age should be calculated as the number of complete years that have passed since the patient’s date of birth [10]. The formula provided in the evidence simply subtracts the year of birth from the current year, which is not accurate.

Examples of E4. In the thrombosis_prediction database, both the Examination and Patient tables contain the Diagnosis column. However, question 1175 asks for the “Doctor’s diagnosis,” and neither the question nor the provided external knowledge specifies which table should be referenced.

Representative Errors in Spider 2.0-Snow. We identify annotation errors in 76 out of the 121 examples in Spider 2.0-Snow that have released gold queries. We show representative error examples for each error pattern in Table 3, and discuss the details below:

Examples of E1. We identify two typical errors for the error pattern E1. Firstly, in four examples, we find that Spider 2.0 annotators used BETWEEN TO_TIMESTAMP(start_date) AND TO_TIMESTAMP(end_date). However, TO_TIMESTAMP() casts a date to the timestamp of the start of the day (00:00:00.000) rather than the end (23:59:59.999). Accordingly, the gold queries omit rows on the end date that occur after midnight, resulting in the wrong query results.

Secondly, we find that three annotated queries omit the required restrictions specified in the question. For example, in sf_bq209, the question asks for a “utility patent,” but the annotated query does not include this restriction.

Examples of E2. The annotators did not filter or deduplicate the intermediate results after the JOIN or FLATTEN operations, which resulted in row inflation in four examples.

Examples of E3. We find that two annotated SQL queries involving forward citation calculation use incorrect join conditions. The queries mistakenly join on cited.patent_id = apps.patent_id instead of cited.citation_id = apps.patent_id.⁷

Examples of E4. We identify seven questions with ambiguous output format specifications. Agents must generate SQL queries that replace quotation marks with the empty string (‘’) to pass the evaluation. Otherwise, the evaluator marks all answers incorrect. However, the questions do not state this requirement.

5 ANNOTATION ERRORS BREAK THE BIRD LEADERBOARD

The leaderboards of text-to-SQL benchmarks provide guidance for both text-to-SQL research and practical deployment. Researchers and practitioners use agent rankings to understand technique gaps and choose agents for downstream applications. To understand the impact of annotation errors in text-to-SQL benchmarks on agent performance and rankings, we evaluated all open-source text-to-SQL agents on the BIRD leaderboard whose reported results are reproducible in local environments. Our evaluation demonstrates the unreliability of the BIRD leaderboard.

In this section, we first describe how we corrected a subset of BIRD and then detail the experimental settings. We next present changes in execution accuracy and rankings of all evaluated agents.

⁷In sf_bq128, the annotators used the correct join key.

Annotated user question: Among the drivers who participated in the French Grand Prix, who has the slowest time in the 3rd lap.

Annotated external knowledge: Slowest time refers to Max(time).

Annotated SQL

```
SELECT T1.driverId FROM lapTimes AS T1
INNER JOIN races AS T2 on T1.raceId = T2.raceId
WHERE T2.name = 'French_Grand_Prix' AND T1.lap = 3
ORDER BY T1.time DESC LIMIT 1
```

Issue

1. Ordering by the text “time” column can misrank values lexicographically. (E2)
2. It is unclear whether to output the driver’s name or ID. (E4)
3. There are multiple races named “French Grand Prix”; the question could refer to any single race or all of them combined. (E4)

(a) Example 985 of BIRD Dev. The query lacks data type conversion (E2), and the output column is ambiguous (E4).

Fixed user question: Among all the drivers who have participated in any year of the French Grand Prix throughout its history, which driver recorded the slowest time on the 3rd lap in any of these races? Please specify the driver id.

Modification to the database

```
INSERT INTO "drivers" ("driverId", "driverRef", "forename", "surname", "url")
VALUES (1111, "test", "test", "test", "test");
INSERT INTO "lapTimes" ("raceId", "driverId", "lap", "time", "milliseconds")
VALUES (25, 1111, 3, "11:00.365", 660365);
```

Fixed SQL

```
SELECT T1.driverId FROM lapTimes AS T1
INNER JOIN races AS T2 on T1.raceId = T2.raceId
WHERE T2.name = 'French_Grand_Prix' AND T1.lap = 3
ORDER BY (CAST(substr(lapTimes.time, 1, instr(lapTimes
.time, ':')-1) AS INTEGER) * 60 * 1000) + (CAST(
substr(lapTimes.time, instr(lapTimes.time, ':')
+1) AS REAL) * 1000) DESC LIMIT 1;
```

(b) Corrected example 985 in BIRD Dev by correcting the user question and the ground-truth SQL query. We also update the database to prevent incorrect queries from producing the same results.

Figure 5: Original and corrected example 985 in BIRD Dev.

We analyze the causes of these performance variations to understand how annotation errors affect text-to-SQL agents. Finally, we analyze two rank correlations: (i) between the agent rankings on the original Dev subset and the full Dev set, and (ii) between the agent rankings on the original and corrected Dev subsets.

5.1 Data Settings

Dataset. We randomly sampled 100 examples from the BIRD Dev set. As shown in Table 4, our selection includes 62 simple, 28 moderate, and 10 challenging examples.

Benchmark Correction. We manually corrected the sampled Dev subset based on SAPAR. For each text-to-SQL example, according to the diagnostic report generated by SAR-Agent, we made the following corrections. First, following prior work [4], we corrected all identified issues by revising natural language questions, external knowledge, and ground truth SQL queries to generate a corrected Dev subset. Second, as mentioned in Section 4.4, the value description of the r type column was previously marked as “unuseful.” We

Table 4: We corrected a total of 48 text-to-SQL examples. Each example may involve more than one type of correction.

Statistics	# Examples (total=100)
Difficulty of selected examples	
Easy	62
Moderate	28
Challenging	10
Correction type of Dev subset	
Corrected examples	48
Question correction	19
External knowledge correction	17
SQL correction	41
Schema correction	1
Database modification	6

corrected the schema to fix this issue in the revised Dev subset. Finally, we found examples where incorrect queries produced the same results as the corrected queries due to limitations in the underlying data. To enable more accurate evaluation, we addressed this issue by inserting distinguishing records in the database. We conducted up to three iterations of review and revision to reach consensus with SAR-Agent for each text-to-SQL pair.

To understand our correction process, we use the Problem 985 of BIRD Dev as an example (Figure 5). We found that both the original user query (in natural language) and gold SQL query are incorrect. Therefore, we corrected both the user query and SQL query. Furthermore, we identified that the original data failed to expose errors of incorrectly ranking time. To address this, we inserted a row with the time column set to “11:00.365” into the database.

Statistics of Correction. As shown in Table 4, we corrected 48% of the examples. Specifically, we corrected the questions in 19 examples, addressed external knowledge issues in 17 examples, corrected the SQL queries in 41 examples, and fixed the schema in one example. In addition, we also modified the data in the database for six examples to prevent incorrect SQL logic from being mistakenly considered correct.

5.2 Experimental Settings

Metrics. Following prior work [34], we use Execution Accuracy (EX) and ranking as our metrics. EX determines correctness by comparing execution results between predicted SQL queries and ground truth queries on the database. The ranking metric determines each text-to-SQL agent’s position among all candidates. We further use Spearman’s rank correlation coefficient (r_s) [53] for rank correlation analysis, which measures monotonic relationships between rank variables, ranging from -1 to $+1$.

Text-to-SQL Agents. We evaluated all 16 agents (Contextual-SQL [1], CSC-SQL [52], GenaSQL [19], OpenSearch-SQL [60], OmniSQL [32], CHESS [54], Alpha-SQL [31], GSR [3], RSL-SQL [9], E-SQL [8], TA-SQL [49], DTS-SQL [46], MAC-SQL [57], SFT CodeS-15B [33], DAIL-SQL [20], DIN-SQL [44]) on the BIRD leaderboard that have publicly available execution code and can be reproduced

in local environments.⁸ We show the detailed configurations of these agents in the supplementary material [25].

5.3 Evaluation Results

We executed all 16 agents on both the original and corrected BIRD Dev subsets to evaluate how annotation errors affect agent performance. We will discuss the change of execution accuracy and rankings of all evaluated agents in the following paragraphs.

Execution Accuracy (EX). As shown in Figure 6, the relative changes in EX across the agents range from -7% to 31% . Specifically, 12 out of 16 agents exhibit improved EX after correction. Among these, CHESS demonstrates the most significant increase, with EX rising from 62% to 81% , a 30.6% relative improvement. In contrast, four agents experience a performance decline. SFT CodeS-15B shows the most significant decrease, with EX dropping from 56% to 52% , corresponding to a 7.1% relative decline. We provide a detailed analysis for the performance change of CHESS and SFT CodeS-15B in Section 5.4.

We then analyze the factors responsible for the changes in EX. We present the number of examples for which correctness changed, alongside information on whether annotation corrections were applied in Figure 7. Across all 16 agents, annotation corrections change their correctness on an average of 16.4 examples from False to True and 8.0 from True to False. Among the 16 agents, CHESS shows the largest number of False to True updates (26), while DTS-SQL shows the fewest (7). In contrast, for True to False changes, Contextual-SQL and OpenSearch-SQL tie for the most (14), and MAC-SQL has the fewest (3).

Rankings. We compare agent rankings before and after correction and find significant changes in their relative positions. As shown in Figure 8, CHESS and GenaSQL, which achieve substantial relative performance improvements of 30.6% and 26.6% respectively, elevate from middle-ranked positions (7th and 6th place) to tie for first place with 81% EX. MAC-SQL shows the most dramatic ranking shift, moving 9 positions from 14th to 5th place due to the second-highest relative performance gain of 29.3% . In contrast, GSR experiences the sharpest decline, falling 9 positions from 4th to 13th place. Despite achieving a 3% relative improvement in performance, GSR’s marginal gain becomes insignificant compared to other improving agents, which average 21% relative improvement.

All agents exhibiting a decline in EX experience ranking drops. Previous top-ranking agents, Contextual-SQL and OpenSearch-SQL, suffer notable repositioning due to performance degradation. Contextual-SQL falls 6 positions from 1st to 7th place, while OpenSearch-SQL drops 9 positions from 2nd to 11th place. In addition, DTS-SQL and SFT CodeS-15B also decline by 6 and 1 positions, respectively.

Across all 16 agents, ranking changes average 5 positions, with individual shifts ranging from -9 to $+9$ positions. These results demonstrate that rankings of methods are highly sensitive to the annotation errors in the benchmark. High error rates present in the benchmark can significantly undermine the reliability of a leaderboard, which fails to accurately represent the relative performance of different agents. Consider a practitioner seeking to select the best

⁸We chose agents based on leaderboard rankings as of August 20, 2025.




Figure 6: Execution accuracy of agents on original and corrected BIRD Dev subsets. The agents’ performance changes range from -7% to 31% in relative terms.




Figure 7: Number of examples with correctness changes between original and corrected BIRD Dev subsets.

text-to-SQL agent for a downstream application based on the BIRD leaderboard. The flawed leaderboard can mislead this decision and lead to the selection of a suboptimal agent.

5.4 Analysis of Performance Changes

In this subsection, we analyze why agents experience varying degrees of performance change from annotation corrections by investigating CHES [54] and SFT CodeS-15B [33], the agents showing the largest improvement and the largest degradation.

Why does the EX of CHES increase? As shown in Table 5, we identify 12 examples where correcting errors only in Q (the gold query) changes the evaluation status from False to True for CHES. We analyze the CHES-generated queries for these 12 examples in the original Dev set. We find that 91.7% of generations match the revised ground truth, indicating that the original False labels are false negatives. Across these 12 examples, we find a common error pattern: in 5/12 (41.7%) of the examples (310, 416, 605, 1286, 1302), human annotators omitted the required DISTINCT keyword within the COUNT function, leading to incorrect False classifications in the original evaluation. In addition, among the examples that require revising both \mathcal{T} (the question or external knowledge) and Q (the gold query) to resolve ambiguities and semantic inconsistencies, we find that 14 CHES-generated queries are reclassified as True.

Why does the EX of SFT CodeS-15B decrease? As shown in Table 6, after correcting Q (the gold query), six examples flip from False to True, whereas three flip from True to False. Of the latter three, two (1286, 1302) change due to the agent committed the




Figure 8: Agent ranking changes from original to corrected BIRD Dev subsets. Ranking shifts range from -9 to +9 positions, with an average change of 5 positions.

same error as the human annotator (missing DISTINCT keyword issue in COUNT function). For examples requiring revisions to both \mathcal{T} (the question or external knowledge) and Q , only three switch from False to True, whereas ten switch from True to False.

5.5 Analysis of Ranking Correlations

To assess the impact of annotation errors on rankings, we conducted two correlation analyses using Spearman’s rank correlation

Table 5: Examples with correctness changes for CHESS categorized by annotation correction type.

Category		Question IDs
Corrected \mathcal{T}	T \rightarrow F	255
Corrected Q	F \rightarrow T	310, 416, 442, 484, 605, 610, 646, 888, 928, 1200, 1286, 1302
	T \rightarrow F	42
Corrected \mathcal{T} & Q	F \rightarrow T	180, 305, 406, 428, 602, 620, 648, 772, 855, 970, 987, 1004, 1271, 1280
	T \rightarrow F	829, 1173

Table 6: Examples with correctness changes for SFT CodeS-15B categorized by annotation correction type.

Category		Question IDs
Corrected Q	F \rightarrow T	442, 605, 610, 646, 888, 928
	T \rightarrow F	1200, 1286, 1302
Corrected \mathcal{T} & Q	F \rightarrow T	406, 855, 1271
	T \rightarrow F	305, 620, 648, 772, 829, 846, 935, 987, 1173, 1280

coefficient [53]. We first examined the correlation between EX rankings on the original Dev subset and the BIRD leaderboard (Dev). As shown in Figure 9a, the coefficient r_s is 0.85 with p -value less than 0.001. This strong positive correlation confirms that the relative performance characteristics of agents on the full BIRD Dev set are preserved in our randomly sampled examples.

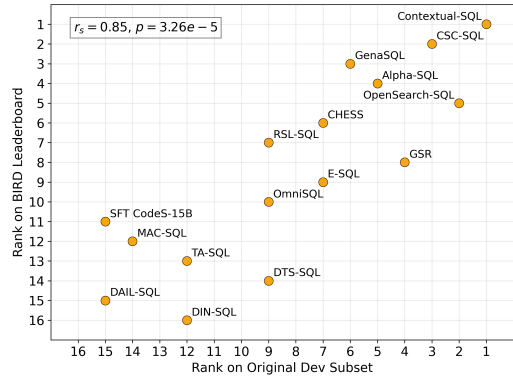
We also compared rankings on the original Dev subset and our corrected Dev subset, where the only factor that differs is whether annotation errors are corrected. As shown in Figure 9b, the coefficient r_s is 0.32 with a p -value of 0.23. This weak correlation indicates that the presence of annotation errors substantially affects ranking outcomes, compromising the reliability of rankings of agents on the original dataset.

6 SAR-AGENT HELPS EXPERTS DETECT 42% MORE INCORRECT EXAMPLES

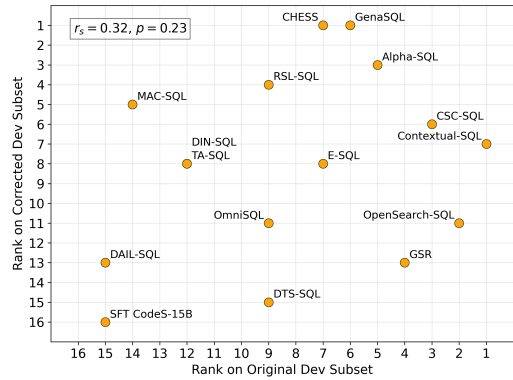
We evaluated the ability of SAR-Agent in detecting text-to-SQL annotation errors. In this section, we first describe the precision of SAR-Agent on BIRD Mini-Dev and Spider 2.0-Snow. We then compare the number of detected error examples by SAR-Agent with those by the Arcwise team. Finally, we present the per-example monetary cost and the number of execution steps of SAR-Agent.

Building upon the corrections by Arcwise, we addressed additional errors identified in BIRD Mini-Dev [4]. To facilitate fair evaluation, we have released two versions of the dataset: (1) Arcwise-Plat-SQL: we only corrected the SQL annotations, preserving original ambiguities in the questions and evidence; (2) Arcwise-Plat: we resolved both incorrect SQL annotations and underlying ambiguities. Both versions are available in our GitHub repository.⁹

⁹https://github.com/uiuc-kang-lab/text_to_sql_benchmarks/tree/main/data



(a) Correlation between EX rankings on the original Dev subset and the full BIRD Dev set. $r_s = 0.85$ indicates a strong positive correlation.



(b) Correlation between EX rankings on the original and corrected Dev subset. $r_s = 0.32$ indicates weak correlation.

Figure 9: Ranking Correlation Analysis using Spearman’s Rank Correlation Coefficient r_s .

6.1 Experimental Settings

Data settings. We analyze SAR-Agent’s performance on BIRD Mini-Dev and Spider 2.0-Snow based on the generated diagnostic reports in error detection experiments (Section 4).

Metrics. We measure the precision of SAR-Agent on both BIRD Mini-Dev and Spider 2.0-Snow. In addition, the Arcwise team reports 161 error examples in BIRD Mini-Dev through human audit [4]. We analyze the SAR-Agent hit rate on these examples, defined as the percentage of Arcwise detections that SAR-Agent also detects. Finally, we measure per-example monetary cost (USD) and the number of execution steps for SAR-Agent.

Evaluation methods. Instead of outputting only a binary YES/NO decision when judging annotations, we collected the agent’s diagnostic report and manually verified the validity of its listed error reasons. We count an agent’s detection as successful for a given text-to-SQL example if manual review confirms at least one of the error reasons it reports.

We further compared its diagnostic reports against the results of the Arcwise team’s manual examination. Specifically, we counted:

Table 7: Performance of SAR-Agent. On BIRD Mini-Dev, SAR-Agent achieves 83% precision with an average cost of \$0.44 and 5.1 steps. On Spider 2.0-Snow, it achieves 89% precision with an average cost of \$1.11 and 7.6 steps.

Benchmarks	BIRD Mini-Dev	Spider 2.0-Snow
Precision (%)	83	89
Average Cost (\$)	0.44	1.11
Average Steps	5.1	7.6




Figure 10: Number of detected error examples in BIRD Mini-Dev by SAR-Agent (after manual verification) and the Arcwise team. Of the 161 noisy examples identified by Arcwise, SAR-Agent achieves a hit rate of 85.7%. Additionally, SAR-Agent identifies 41.6% more error examples than Arcwise.


(1) examples flagged as incorrect by both SAR-Agent and the Arcwise team, (2) examples flagged only by SAR-Agent, and (3) examples flagged only by the Arcwise team.

6.2 Evaluation Results


Precision of SAR-Agent. We evaluated SAR-Agent on BIRD Mini-Dev and Spider 2.0-Snow. We report the precision of SAR-Agent for both benchmarks in Table 7. Among 498 examples in BIRD Mini-Dev, SAR-Agent flags 274 as incorrect. Our manual verification confirms 228 as true positive, yielding a precision of 83%. On Spider 2.0-Snow, SAR-Agent flags 79 examples as incorrect. After manual verification, we find the provided error reasons are valid in 70 cases, yielding a precision of 89%.

Comparison of SAR-Agent with Arcwise SQL Experts. We compare SAR-Agent against the SQL experts from the Arcwise team in Figure 10. After manual verification, we find that SAR-Agent successfully detects 138 of the 161 error examples reported by the Arcwise team, achieving an 85.7% hit rate. Additionally, we compare the total number of error examples detected by SAR-Agent and the Arcwise team. SAR-Agent detects 41.6% more error examples than the Arcwise team.


Cost and execution steps of SAR-Agent. We report SAR-Agent’s monetary cost and the number of execution steps per example in Table 7. On BIRD Mini-Dev, SAR-Agent spends \$0.44 per example with an average of 5.1 steps. On Spider 2.0-Snow, which features more complex queries and schemas, SAR-Agent requires more annotation-diagnosis steps, increasing the per-example cost to \$1.11 and the average number of steps to 7.6.



(a) Correlation between EX rankings on the corrected Dev subset with and without database modification. $r_s = 0.95$ indicates a strong positive correlation.



(b) Correlation between EX rankings on the fully corrected Dev subset and the SQL-only-corrected Dev subset by SAR-Agent. $r_s = 0.73$ indicates strong correlation.



(c) Correlation between EX rankings on the fully corrected Dev subset and the SQL-only-corrected Dev subset by humans. $r_s = 0.59$ indicates moderate correlation.

Figure 11: Rank correlation between the fully corrected Dev subset and corrected Dev subsets across settings.

7 ABLATION STUDY

We conducted three Spearman’s rank correlation analyses to assess how different correction settings (i.e., (i) no database modification,

(ii) SQL-only-corrected Dev subset by SAR-Agent, and (iii) SQL-only-corrected Dev subset by humans) affect agent rankings. For each setting, we compare EX rankings against those on the fully corrected Dev subset (Table 4). We describe the details of each correction setting in the supplementary material [25].

We first compared EX rankings on the corrected Dev subset with and without database modification. As shown in Figure 11a, $r_s = 0.95$ indicates that database modification has minimal impact on the rankings of agents.

We next evaluated SQL-only corrections, considering fixes produced by SAR-Agent and by humans. As shown in Figure 11b and Figure 11c, the correlation between EX rankings on the fully corrected Dev subset and the SQL-only-corrected Dev subset by SAR-Agent ($r_s = 0.73$) is stronger than that between the fully corrected Dev subset and the SQL-only-corrected subset by humans ($r_s = 0.59$). We attribute this gap to residual issues in the natural language input (e.g., ambiguity or domain-knowledge errors). SAR-Agent’s fixes assume these issues are resolved, making its SQL-only corrections closer to the fully corrected Dev subset. In contrast, human-repaired SQL queries align with the question and external knowledge as written, even when they contain issues.

8 RELATED WORK

Text-to-SQL benchmarks. Researchers have proposed a wide range of datasets and benchmarks to evaluate the effectiveness of text-to-SQL approaches. Early research primarily uses single-domain datasets [24, 48, 61]. In recent years, attention has shifted toward cross-domain datasets to better assess models’ generalization across diverse database schemas [29, 34, 63, 65]. The rapid progress in LLMs and text-to-SQL techniques has further driven the development of new benchmarks, which feature more complex SQL queries and database schemas [13, 30, 35, 55, 56]. In this work, we systematically analyze errors in two widely studied text-to-SQL benchmarks, BIRD Mini-Dev [34] and Spider 2.0-Snow [30]. Our analysis reveals error rates exceeding 50% across these benchmarks.

Text-to-SQL methods. Text-to-SQL approaches have a long history, with research evolving significantly in recent years thanks to advances in LLMs. Recent work has focused on developing text-to-SQL agents [1, 3, 8, 9, 17, 19, 20, 31, 44, 46, 49, 52, 54, 57, 60] that decompose the task into several sub-tasks and modules, including schema linking, SQL generation, and SQL refinement. Another major direction is fine-tuning LLMs for text-to-SQL tasks. Early efforts primarily relies on supervised learning [21, 32], whereas recent approaches increasingly leverage reinforcement learning frameworks to strengthen the reasoning capabilities of LLMs [47, 52]. In this work, we re-evaluated all open-source agents on the BIRD leaderboard using both the original and corrected versions of the BIRD Dev subset. The re-evaluation reveals significant changes in execution accuracy and agent rankings, demonstrating that annotation errors undermine the leaderboard’s reliability.

Annotation errors in benchmarks. Many studies have highlighted significant issues with existing benchmarks [2, 4, 14, 27, 36, 39, 45, 59, 62, 64, 68]. In particular, several works have examined the presence of annotation errors in widely-used text-to-SQL benchmarks such as Spider and BIRD [4, 39, 45, 59]. To fix errors in

agentic benchmarks, prior works use the combination of automatic correction and manual verification. For example, in coding and software engineering benchmarks [14, 36, 64], researchers employ automatically generated test cases and oracles to strengthen the rigor of evaluating code generation. Moreover, manual verification and checks are used to ensure the soundness of error corrections [2, 14, 36, 64, 68]. Our work combines agent-powered automation with manual verification to audit and correct text-to-SQL benchmarks. We ran SAR-Agent to automatically detect and diagnose text-to-SQL annotation errors, with every reported error subsequently reviewed by SQL experts. Final adjudication on whether an annotation is incorrect was made by SQL experts. We corrected a sampled BIRD Dev subset based on SAPAR. In SAPAR, human annotators can adopt the SQL revision proposed by SAR-Agent after manual verification. If they disagree, they revise the SQL themselves. Annotators also need to revise the user question and supporting evidence when these contain errors.

LLM-as-a-Judge. LLM-as-a-Judge denotes using LLMs, which are typically aligned with human preferences via reinforcement learning from human feedback (RLHF), to evaluate system performance [67]. As LLMs’ capabilities, particularly reasoning, have improved, LLM-as-a-Judge has been integrated into evaluation workflows [22] across finance [6, 58, 66], law [38, 51], and science [5, 28]. Recent work extends this paradigm to multimodal scenarios with multimodal large language models (MLLMs), termed MLLM-as-a-Judge [11], and to agentic frameworks for evaluating agentic systems, termed Agent-as-a-Judge [69]. Because Text-to-SQL annotation requires not only database expertise but also domain knowledge across fields such as science, sports, and finance, we develop the first AI agent (SAR-Agent) to serve as a co-reviewer for text-to-SQL annotation. We hope SAR-Agent will facilitate the development of high-quality text-to-SQL benchmarks.

9 CONCLUSION

In this work, we quantify annotation error rates for two widely used text-to-SQL benchmarks, BIRD and Spider 2.0-Snow. We then correct a subset of the BIRD Dev set to systematically evaluate how annotation errors affect agent performance and leaderboard rankings. To facilitate error detection, we develop SAR-Agent, which generates a diagnostic report for each text-to-SQL annotation. After SQL experts manually verify each error reason listed in the report, we find error rates of 52.8% and 62.8% in BIRD Mini-Dev and Spider 2.0-Snow, respectively. We introduce SAPAR, which incorporates SAR-Agent into the existing text-to-SQL annotation pipeline, to correct a sampled BIRD Dev set. Re-evaluating all 16 open-source text-to-SQL agents from the BIRD leaderboard, we find relative performance changes from -7% to +31% and ranking shifts from -9 to +9 positions, underscoring the unreliability of current text-to-SQL leaderboards. We advocate that future work use SAR-Agent and SAPAR to develop high-quality text-to-SQL benchmarks.

10 ACKNOWLEDGEMENTS

We are grateful to the CloudLab [18] for providing computing resources for experiments. This research was supported in part by Google and the Open Philanthropy project.

REFERENCES

- [1] Sheshansh Agrawal and Thien Nguyen. 2025. Open-Sourcing the Best Local Text-to-SQL System. <https://contextual.ai/blog/open-sourcing-the-best-local-text-to-sql-system/>
- [2] Reem Aleithan, Haoran Xue, Mohammad Mahdi Mohajer, Elijah Nnorom, Gias Uddin, and Song Wang. 2024. SWE-Bench+: Enhanced Coding Benchmark for LLMs. arXiv:2410.06992 [cs.SE] <https://arxiv.org/abs/2410.06992>
- [3] Anonymous. 2025. LLM Prompting for Text2SQL via Gradual SQL Refinement. In *Submitted to ACL Rolling Review - February 2025*. <https://openreview.net/forum?id=ozUJ0ijVTJ> under review.
- [4] Arcwise. 2024. BIRD Mini Dev Annotation Issues. <https://docs.google.com/spreadsheets/d/1IGm9Otruy60ujUnl8AOkepY3qgWdFJHnX7hQGUEcW/edit?gid=0#gid=0>
- [5] Nathan Brake and Thomas Schaaf. 2024. Comparing Two Model Designs for Clinical Note Generation: Is an LLM a Useful Evaluator of Consistency?. In *Findings of the Association for Computational Linguistics: NAACL 2024*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 352–363. <https://doi.org/10.18653/v1/2024.findings-naacl.25>
- [6] Meni Brief, Oded Ovadia, Gil Shenderovitz, Noga Ben Yoash, Rachel Lemberg, and Eitam Sheerit. 2024. Mixing It Up: The Cocktail Effect of Multi-Task Fine-Tuning on LLM Performance – A Case Study in Finance. arXiv:2410.01109 [cs.AI] <https://arxiv.org/abs/2410.01109>
- [7] Ben Bush. 2025. What Is The 90% Rule In F1? A Complete Guide. <https://www.formulaonehistory.com/what-is-the-90-rule-in-f1-a-complete-guide/>
- [8] Hasan Alp Caferoğlu and Özgür Ulusoy. 2024. E-SQL: Direct Schema Linking via Question Enrichment in Text-to-SQL. arXiv:2409.16751 [cs.CL] <https://arxiv.org/abs/2409.16751>
- [9] Zhenbiao Cao, Yuanlei Zheng, Zhihao Fan, Xiaojin Zhang, and Wei Chen. 2024. RSL-SQL: Robust Schema Linking in Text-to-SQL Generation. *arXiv preprint arXiv:2411.00073* (2024).
- [10] CDC. 2025. Age. <https://www.cdc.gov/nchs/hus/sources-definitions/age.htm>
- [11] Dongping Chen, Ruoxi Chen, Shilin Zhang, YINUO Liu, Yaochen Wang, Huichi Zhou, Qihui Zhang, Yao Wan, Pan Zhou, and Lichao Sun. 2024. MLLM-as-a-Judge: Assessing Multimodal LLM-as-a-Judge with Vision-Language Benchmark. arXiv:2402.04788 [cs.CL] <https://arxiv.org/abs/2402.04788>
- [12] Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. BGE M3-Embedding: Multi-Lingual, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation. arXiv:2402.03216 [cs.CL] <https://arxiv.org/abs/2402.03216>
- [13] Peter Baile Chen, Fabian Wenz, Yi Zhang, Devin Yang, Justin Choi, Nesime Tatbul, Michael Cafarella, Çağatay Demiralp, and Michael Stonebraker. 2025. BEAVER: An Enterprise Benchmark for Text-to-SQL. arXiv:2409.02038 [cs.CL] <https://arxiv.org/abs/2409.02038>
- [14] Zhi Chen and Lingxiao Jiang. 2025. Evaluating Software Development Agents: Patch Patterns, Code Quality, and Issue Complexity in Real-World GitHub Scenarios. In *2025 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 657–668. <https://doi.org/10.1109/SANER64311.2025.00068>
- [15] Wikipedia contributors. 2025. K-12. <https://en.wikipedia.org/wiki/K-12>
- [16] DeepSeek-AI. 2024. DeepSeek LLM: Scaling Open-Source Language Models with Longtermism. *arXiv preprint arXiv:2401.02954* (2024). <https://github.com/deepseek-ai/DeepSeek-LLM>
- [17] Minghang Deng, Ashwin Ramachandran, Canwen Xu, Lanxiang Hu, Zhewei Yao, Anupam Datta, and Hao Zhang. 2025. ReFoRCE: A Text-to-SQL Agent with Self-Refinement, Consensus Enforcement, and Column Exploration. arXiv:2502.00675 [cs.CL] <https://arxiv.org/abs/2502.00675>
- [18] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar, and Prabhodh Mishra. 2019. The Design and Operation of CloudLab. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*. USENIX Association, Renton, WA, 1–14. <https://www.usenix.org/conference/atc19/presentation/duplyakin>
- [19] Yusuf Denizay Dönder, Derek Hommel, Andrea W Wen-Yi, David Mimmo, and Unso Eun Seo Jo. 2025. Cheaper, Better, Faster, Stronger: Robust Text-to-SQL without Chain-of-Thought or Fine-Tuning. arXiv:2505.14174 [cs.CL] <https://arxiv.org/abs/2505.14174>
- [20] Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. 2024. Text-to-SQL Empowered by Large Language Models: A Benchmark Evaluation. *Proc. VLDB Endow.* 17, 5 (Jan. 2024), 1132–1145. <https://doi.org/10.14778/3641204.3641221>
- [21] Yingqi Gao, Yifu Liu, Xiaoxia Li, Xiaorong Shi, Yin Zhu, Yiming Wang, Shiqi Li, Wei Li, Yuntao Hong, Zhiling Luo, Jinyang Gao, Liyu Mou, and Yu Li. 2025. A Preview of XiYan-SQL: A Multi-Generator Ensemble Framework for Text-to-SQL. arXiv:2411.08599 [cs.AI] <https://arxiv.org/abs/2411.08599>
- [22] Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Yuanzhuo Wang, Wen Gao, Lionel Ni, and Jian Guo. 2025. A Survey on LLM-as-a-Judge. arXiv:2411.15594 [cs.CL] <https://arxiv.org/abs/2411.15594>
- [23] IBM. 2025. What is exploratory data analysis (EDA)? (2025). <https://www.ibm.com/think/topics/exploratory-data-analysis>
- [24] Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. Learning a Neural Semantic Parser from User Feedback. arXiv:1704.08760 [cs.CL] <https://arxiv.org/abs/1704.08760>
- [25] Tengjun Jin, Yoojin Choi, Yuxuan Zhu, and Daniel Kang. 2025. Pervasive Annotation Errors Break Text-to-SQL Benchmarks and Leaderboards (Supplementary Material). https://github.com/uiuc-kang-lab/text_to_sql_benchmarks/blob/main/materials/supplementary_material.pdf
- [26] Daniel Kang, Alex Derhacopian, Kaoru Tsuji, Trevor Hebert, Peter Bailis, Tadashi Fukami, Tatsunori Hashimoto, Yi Sun, and Matei Zaharia. [n.d.]. Exploiting Proximity Search and Easy Examples to Select Rare Events. *NeurIPS Data-Centric AI Workshop 2021* ([n. d.]). <https://par.nsf.gov/biblio/10316595>
- [27] Sayash Kapoor, Benedikt Stroebel, Zachary S. Siegel, Nitya Nadgir, and Arvind Narayanan. 2024. AI Agents That Matter. arXiv:2407.01502 [cs.LG] <https://arxiv.org/abs/2407.01502>
- [28] Jack Krolik, Herprit Mahal, Feroz Ahmad, Gaurav Trivedi, and Bahador Saket. 2024. Towards Leveraging Large Language Models for Automated Medical Q&A Evaluation. arXiv:2409.01941 [cs.CL] <https://arxiv.org/abs/2409.01941>
- [29] Chia-Hsuan Lee, Oleksandr Polozov, and Matthew Richardson. 2021. KaggleDBQA: Realistic Evaluation of Text-to-SQL Parsers. arXiv:2106.11455 [cs.CL] <https://arxiv.org/abs/2106.11455>
- [30] Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, Victor Zhong, Caiming Xiong, Ruoxi Sun, Qian Liu, Sida Wang, and Tao Yu. 2024. Spider 2.0: Evaluating Language Models on Real-World Enterprise Text-to-SQL Workflows. arXiv:2411.07763 [cs.CL] <https://arxiv.org/abs/2411.07763>
- [31] Boyan Li, Jiayi Zhang, Ju Fan, Yanwei Xu, Chong Chen, Nan Tang, and Yuyu Luo. 2025. Alpha-SQL: Zero-Shot Text-to-SQL using Monte Carlo Tree Search. arXiv:2502.17248 [cs.DB] <https://arxiv.org/abs/2502.17248>
- [32] Haoyang Li, Shang Wu, Xiaokang Zhang, Xinmei Huang, Jing Zhang, Fuxin Jiang, Shuai Wang, Tieying Zhang, Jianjun Chen, Rui Shi, Hong Chen, and Cuiping Li. 2025. OmniSQL: Synthesizing High-quality Text-to-SQL Data at Scale. arXiv:2503.02240 [cs.CL] <https://arxiv.org/abs/2503.02240>
- [33] Haoyang Li, Jing Zhang, Hanbing Liu, Ju Fan, Xiaokang Zhang, Jun Zhu, Renjie Wei, Hongyan Pan, Cuiping Li, and Hong Chen. 2024. CodeS: Towards Building Open-source Language Models for Text-to-SQL. arXiv:2402.16347 [cs.CL] <https://arxiv.org/abs/2402.16347>
- [34] Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin C.C. Chang, Fei Huang, Reynold Cheng, and Yongbin Li. 2024. Can LLM already serve as a database interface? a big bench for large-scale database grounded text-to-SQLs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 1835, 28 pages.
- [35] Jinyang Li, Xiaolong Li, Ge Qu, Per Jacobsson, Bowen Qin, Binyuan Hui, Shuzheng Si, Nan Huo, Xiaohan Xu, Yue Zhang, Ziwei Tang, Yuanshuai Li, Florensia Widjaja, Xintong Zhu, Feige Zhou, Yongfeng Huang, Yannis Papakonstantinou, Fatma Ozcan, Chenhao Ma, and Reynold Cheng. 2025. SWE-SQL: Illuminating LLM Pathways to Solve User SQL Issues in Real-World Applications. arXiv:2506.18951 [cs.DB] <https://arxiv.org/abs/2506.18951>
- [36] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. arXiv:2305.01210 [cs.SE] <https://arxiv.org/abs/2305.01210>
- [37] Shu Liu, Sumanth Hegde, Shiyi Cao, Alan Zhu, Dacheng Li, Tyler Griggs, Eric Tang, Akshay Malik, Kourosh Hakhmaneshi, Richard Liaw, Philipp Moritz, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. 2025. SkyRL-SQL: Matching GPT-4o and o4-mini on Text2SQL with Multi-Turn RL. Notion Blog.
- [38] Shengjie Ma, Qi Chu, Jiaxin Mao, Xuhui Jiang, Haozhe Duan, and Chong Chen. 2025. Leveraging Large Language Models for Relevance Judgments in Legal Case Retrieval. arXiv:2403.18405 [cs.AI] <https://arxiv.org/abs/2403.18405>
- [39] NL2SQL-Empirical. [n.d.]. The annotation issues in the dev set of BIRD Benchmark. <https://github.com/AlibabaResearch/DAMO-ConvAI/issues/39#issuecomment-2487904959>
- [40] OpenAI. 2024. GPT-4o System Card. arXiv:2410.21276 [cs.CL] <https://arxiv.org/abs/2410.21276>
- [41] OpenAI. 2024. New embedding models and API updates. <https://openai.com/index/new-embedding-models-and-api-updates/>
- [42] OpenAI. 2025. OpenAI o3 and o4-mini System Card. (2025). <https://cdn.openai.com/pdf/2221c875-02dc-4789-800b-e7758f3722c1/o3-and-o4-mini-system-card.pdf>
- [43] Mohammadreza Pourreza, Hailong Li, Ruoxi Sun, Yeounoh Chung, Shayan Talaei, Gaurav Tarlok Kakkar, Yu Gan, Amin Saberi, Fatma Ozcan, and Sercan O. Arık. 2024. CHASE-SQL: Multi-Path Reasoning and Preference Optimized Candidate Selection in Text-to-SQL. arXiv:2410.01943 [cs.LG] <https://arxiv.org/abs/2410.01943>

- [44] Mohammadreza Pourreza and Davood Rafiei. 2023. DIN-SQL: Decomposed In-Context Learning of Text-to-SQL with Self-Correction. arXiv:2304.11015 [cs.CL] <https://arxiv.org/abs/2304.11015>
- [45] Mohammadreza Pourreza and Davood Rafiei. 2023. Evaluating Cross-Domain Text-to-SQL Models and Benchmarks. arXiv:2310.18538 [cs.CL] <https://arxiv.org/abs/2310.18538>
- [46] Mohammadreza Pourreza and Davood Rafiei. 2024. DTS-SQL: Decomposed Text-to-SQL with Small Large Language Models. arXiv:2402.01117 [cs.CL] <https://arxiv.org/abs/2402.01117>
- [47] Mohammadreza Pourreza, Shayan Talaei, Ruoxi Sun, Xingchen Wan, Hailong Li, Azalia Mirhoseini, Amin Saberi, and Sercan O. Arik. 2025. Reasoning-SQL: Reinforcement Learning with SQL Tailored Partial Rewards for Reasoning-Enhanced Text-to-SQL. arXiv:2503.23157 [cs.LG] <https://arxiv.org/abs/2503.23157>
- [48] P. J. Price. 1990. Evaluation of Spoken Language Systems: the ATIS Domain. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*. <https://aclanthology.org/H90-1020/>
- [49] Ge Qu, Jinyang Li, Bowen Li, Bowen Qin, Nan Huo, Chenhao Ma, and Reynold Cheng. 2024. Before Generation, Align it! A Novel and Effective Strategy for Mitigating Hallucinations in Text-to-SQL Generation. In *Findings of the Association for Computational Linguistics ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand and virtual meeting, 5456–5471. <https://doi.org/10.18653/v1/2024.findings-acl.324>
- [50] Qwen. 2025. Qwen2.5 Technical Report. arXiv:2412.15115 [cs.CL] <https://arxiv.org/abs/2412.15115>
- [51] Cheol Ryu, Seolhwa Lee, Subeeng Pang, Chanyeol Choi, Hojun Choi, Myeonggee Min, and Jy-Yong Sohn. 2023. Retrieval-based Evaluation for LLMs: A Case Study in Korean Legal QA. In *Proceedings of the Natural Language Processing Workshop 2023*, Daniel Preotiu-Pietro, Catalina Goanta, Ilias Chalkidis, Leslie Barrett, Gerasimos Spanakis, and Nikolaos Aletras (Eds.). Association for Computational Linguistics, Singapore, 132–137. <https://doi.org/10.18653/v1/2023.nllp-1.13>
- [52] Lei Sheng and Shuai-Shuai Xu. 2025. CSC-SQL: Corrective Self-Consistency in Text-to-SQL via Reinforcement Learning. arXiv:2505.13271 [cs.CL] <https://arxiv.org/abs/2505.13271>
- [53] Charles Spearman. 1904. The Proof and Measurement of Association between Two Things. *The American Journal of Psychology* 15, 1 (1904), 72–101. <http://www.jstor.org/stable/1412159>
- [54] Shayan Talaei, Mohammadreza Pourreza, Yu-Chen Chang, Azalia Mirhoseini, and Amin Saberi. 2024. CHESS: Contextual Harnessing for Efficient SQL Synthesis. arXiv:2405.16755 [cs.LG] <https://arxiv.org/abs/2405.16755>
- [55] BIRD Team. 2024. LiveSQLBench: A Dynamic and Contamination-Free Benchmark for Evaluating LLMs on Real-World Text-to-SQL Tasks. <https://github.com/bird-bench/livesqlbench>. Accessed: 2025-05-22.
- [56] BIRD Team. 2025. BIRD-Interact: Re-imagine Text-to-SQL Evaluation via Lens of Dynamic Interactions. <https://github.com/bird-bench/BIRD-Interact>. Accessed: 2025-06-01.
- [57] Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, LinZheng Chai, Zhao Yan, Qian-Wen Zhang, Di Yin, Xing Sun, and Zhoujun Li. 2025. MAC-SQL: A Multi-Agent Collaborative Framework for Text-to-SQL. arXiv:2312.11242 [cs.CL] <https://arxiv.org/abs/2312.11242>
- [58] Saizhuo Wang, Hang Yuan, Lionel M. Ni, and Jian Guo. 2024. QuantAgent: Seeking Holy Grail in Trading by Self-Improving Large Language Model. arXiv:2402.03755 [cs.AI] <https://arxiv.org/abs/2402.03755>
- [59] Niklas Wretblad, Fredrik Gordh Riseby, Rahul Biswas, Amin Ahmadi, and Oskar Holmström. 2024. Understanding the Effects of Noise in Text-to-SQL: An Examination of the BIRD-Bench Benchmark. arXiv:2402.12243 [cs.CL] <https://arxiv.org/abs/2402.12243>
- [60] Xiangjin Xie, Guangwei Xu, Lingyan Zhao, and Ruijie Guo. 2025. OpenSearch-SQL: Enhancing Text-to-SQL with Dynamic Few-shot and Consistency Alignment. arXiv:2502.14913 [cs.CL] <https://arxiv.org/abs/2502.14913>
- [61] Navid Yaghmazadeh, Yuepeng Wang, Isil Dillig, and Thomas Dillig. 2017. SQLizer: query synthesis from natural language. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 63 (Oct. 2017), 26 pages. <https://doi.org/10.1145/3133887>
- [62] Ke Yang, Yao Liu, Sapana Chaudhary, Rasool Fakoor, Pratik Chaudhari, George Karypis, and Huzefa Rangwala. 2025. AgentOccam: A Simple Yet Strong Baseline for LLM-Based Web Agents. arXiv:2410.13825 [cs.AI] <https://arxiv.org/abs/2410.13825>
- [63] Semih Yavuz, Izzeddin Gur, Yu Su, and Xifeng Yan. 2018. What It Takes to Achieve 100% Condition Accuracy on WikiSQL. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 1702–1711. <https://doi.org/10.18653/v1/D18-1197>
- [64] Boxi Yu, Yuxuan Zhu, Pinjia He, and Daniel Kang. 2025. UTBoost: Rigorous Evaluation of Coding Agents on SWE-Bench. arXiv:2506.09289 [cs.SE] <https://arxiv.org/abs/2506.09289>
- [65] 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. 2019. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. arXiv:1809.08887 [cs.CL] <https://arxiv.org/abs/1809.08887>
- [66] Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yupeng Cao, Zhi Chen, Jordan W. Suchow, Rong Liu, Zhenyu Cui, Zhaozhuo Xu, Denghui Zhang, Koduvayur Subbalakshmi, Guojun Xiong, Yueru He, Jimin Huang, Dong Li, and Qianqian Xie. 2024. FinCon: A Synthesized LLM Multi-Agent System with Conceptual Verbal Reinforcement for Enhanced Financial Decision Making. arXiv:2407.06567 [cs.CL] <https://arxiv.org/abs/2407.06567>
- [67] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. arXiv:2306.05685 [cs.CL] <https://arxiv.org/abs/2306.05685>
- [68] Yuxuan Zhu, Tengjun Jin, Yada Pruksachatkun, Andy Zhang, Shu Liu, Sasha Cui, Sayash Kapoor, Shayne Longpre, Kevin Meng, Rebecca Weiss, Fazl Barez, Rahul Gupta, Jwala Dhamala, Jacob Merizian, Mario Giulianelli, Harry Cop-pock, Cozmin Ududec, Jasjeet Sekhon, Jacob Steinhardt, Antony Kellerman, Sarah Schwettmann, Matei Zaharia, Ion Stoica, Percy Liang, and Daniel Kang. 2025. Establishing Best Practices for Building Rigorous Agentic Benchmarks. arXiv:2507.02825 [cs.AI] <https://arxiv.org/abs/2507.02825>
- [69] Mingchen Zhuge, Changsheng Zhao, Dylan Ashley, Wenyi Wang, Dmitrii Khizbullin, Yunyang Xiong, Zechun Liu, Ernie Chang, Raghuraman Krishnamoorthi, Yuandong Tian, Yangyang Shi, Vikas Chandra, and Jürgen Schmidhuber. 2024. Agent-as-a-Judge: Evaluate Agents with Agents. arXiv:2410.10934 [cs.AI] <https://arxiv.org/abs/2410.10934>