

# 000 SQLGOVERNOR: AN LLM-POWERED SQL TOOLKIT 001 002 FOR REAL WORLD APPLICATION 003 004

005 **Anonymous authors**

006 Paper under double-blind review

## 007 008 ABSTRACT 009

011 SQL queries in real-world analytical environments—whether written by humans  
012 or generated automatically—often suffer from syntax errors, inefficiency, or se-  
013 mantic misalignment, especially in complex OLAP scenarios. To address these  
014 challenges, we propose SQLGOVERNOR, an LLM-powered SQL toolkit that uni-  
015 fies multiple functionalities—including syntax correction, query rewriting, query  
016 modification, and consistency verification—within a structured framework en-  
017 hanced by knowledge management. SQLGOVERNOR introduces a fragment-wise  
018 processing strategy to enable fine-grained rewriting and localized error correction,  
019 significantly reducing the cognitive load on the LLM. It further incorporates a hy-  
020 brid self-learning mechanism guided by expert feedback, allowing the system to  
021 continuously improve through DBMS output analysis and rule validation. Ex-  
022 periments on benchmarks such as BIRD and BIRD-CRITIC, as well as industrial  
023 datasets, show that SQLGOVERNOR consistently boosts the performance of base  
024 models by up to 10%, while minimizing reliance on manual expertise. Deployed  
025 in production environments, SQLGOVERNOR demonstrates strong practical util-  
026 ity and effective performance.

## 027 1 INTRODUCTION

030 In real-world analytical applications, Structured Query Language (SQL) remains the primary inter-  
031 face for interacting with relational databases. Despite its maturity and widespread adoption, crafting  
032 accurate, efficient, and semantically aligned SQL queries—especially in complex analytical (OLAP)  
033 scenarios—remains a challenging task for both novice and experienced users alike.

034 OLAP workloads are central to modern business intelligence, reporting, and decision-making sys-  
035 tems Codd (1993); Thomsen (2002). Even minor inefficiencies or ambiguities can lead to significant  
036 performance degradation, incorrect insights, or increased development overhead Vassiliadis & Sellis  
037 (1999); Pedersen & Jensen (2001). SQL queries in OLAP settings typically exhibit three key charac-  
038 teristics. First, they perform multi-dimensional analysis using advanced operations such as roll-up  
039 and drill-down Ceci et al. (2015), resulting in highly structured and deeply nested query forms.  
040 Second, these queries operate on large-scale enterprise data Chen et al. (2012), which increases  
041 computational costs and run-time unpredictability. Third, many OLAP queries are executed repeat-  
042 edly—such as daily or weekly reports—making even small inefficiencies costly over time Zhan et al.  
043 (2019).

044 The repetitive and high-stakes nature of OLAP queries amplifies the need for robust, automated, and  
045 adaptive SQL post-processing solutions. Given these challenges, many tools have been developed  
046 to assist users in crafting better SQL queries, including syntax correction, query rewriting, and  
047 semantic refinement Cen et al. (2024); Chen et al. (2023); Liu & Mozafari (2024); Li et al. (2024c).  
048 While large language models (LLMs) have shown great promise in translating natural language  
049 questions into SQL, their applicability across the broader spectrum of SQL-related tasks remains  
050 underexplored. Moreover, in industrial practice, many users lack deep database expertise and often  
051 produce poorly written queries that are inefficient or semantically inaccurate, further exacerbating  
052 system performance and reliability issues Banisharif et al. (2022).

053 *C1: Productivity bottleneck from a fragmented ecosystem.* Existing SQL tools offer isolated func-  
054 tionalities, lacking a unified framework for tasks like syntax correction, semantic refinement, and

query rewriting Cen et al. (2024); Chen et al. (2023); Li et al. (2024c). This fragmentation creates a high barrier to entry for non-experts and increases manual effort by 30-40% for experienced practitioners due to context switching and compatibility issues team (2023a).

*C2: Lack of advanced techniques tailored for OLAP.* Most existing tools target general-purpose or simpler workloads like OLTP, NL2SQL, or offline optimization Chen et al. (2023); Wang et al. (2022); Cen et al. (2024); Yi et al. (2025). OLAP queries are complex, long-running, and require a careful balance between effectiveness and computational cost. Lightweight methods often fail to capture this complexity, while computationally intensive ones risk increasing end-to-end execution time Zhan et al. (2019).

*C3: High operational cost in an expert-centric knowledge lifecycle.* Correcting and rewriting queries demands deep domain and database expertise, which is difficult to capture with conventional data-driven methods. This reliance on expert teams increases labor costs by 25-35% Gartner (2022). Furthermore, maintaining and updating this expert knowledge is time-consuming, limiting scalability and adoption.

In summary, the current landscape lacks a comprehensive and practical SQL toolkit, which utilize evolving knowledge with fewer human efforts.

To address *C1*, we propose an LLM-powered SQL toolkit that unifies multiple functionalities within a structured framework enhanced by knowledge management. Users can either select individual tools for specific tasks or use an end-to-end pipeline that orchestrates multiple tools in a coordinated, use-case-driven manner. By consolidating diverse functionalities into a single platform, our approach eliminates the fragmentation in existing SQL tool-chains, significantly reducing deployment overhead, manual effort, and the barrier to entry.

To address *C2*, we adopt a dual approach that applies validated rules as guidance when appropriate, while permitting the LLM autonomous operation otherwise. Additionally, for particularly long and structurally complex queries, we propose a “fragment processing” strategy to reduce the chance of LLM hallucinations and lower the cost of using the LLM.

To address *C3*, we propose an expert-guided iterative self-learning mechanism to maintain a dynamic knowledge base for SQL tasks. The LLM agent analyzes DBMS outputs to generate new rules, identifying unseen error types from failed SQL and discovering rewriting strategies for inefficient queries. These rules are periodically verified by experts and integrated into the knowledge base for continuous improvement.

The main contributions of this paper can be summarized as follows:

1. **Unified Framework:** To the best of our knowledge, SQLGOVERNOR is the first comprehensive LLM-based SQL toolkit with a knowledge management module. It provides four core functionalities powered by a hybrid self-learning mechanism, thereby improving both user productivity and SQL quality.
2. **Fragment Processing:** We propose a fragment-wise processing strategy to address the complexity and length of OLAP queries. By localizing error detection and rewriting within individual fragments, including subqueries and CTEs, our approach enhances precision and reduces the cognitive burden on LLMs.
3. **Hybrid Self-Learning:** We introduce an expert-guided hybrid self-learning framework that enables SQLGOVERNOR to automatically extract common pattern from execution outputs, generate and validate new knowledge with minimal expert intervention, leading to continuous performance improvement.
4. **Proven Effectiveness:** Extensive experiments on academic benchmarks and real-world industrial datasets demonstrate that SQLGOVERNOR consistently improves the performance of base models by up to 10% in key metrics. Deployed in production environments, SQLGOVERNOR indicates strong practical utility and effective performance.

## 2 FRAMEWORK DESIGN

As illustrated in Figure 1, the architecture of SQLGOVERNOR comprises four specialized tools, each tailored to address a specific category of SQL-related tasks. These tools are supported by

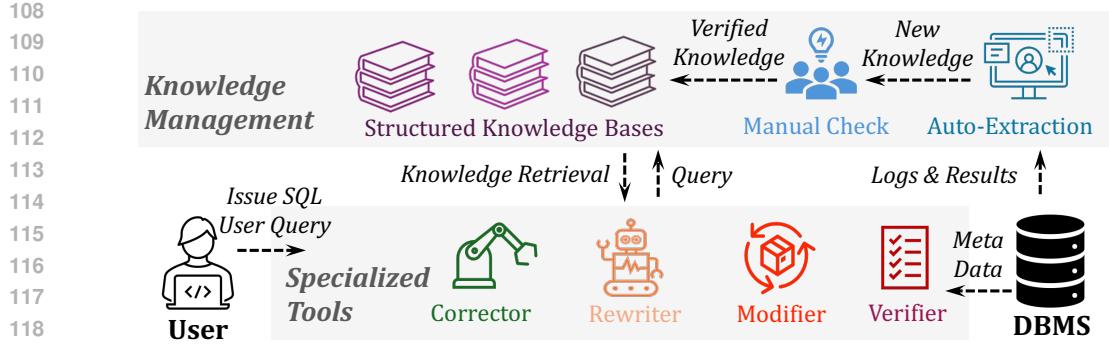


Figure 1: SQLGOVERNOR integrates four specialized SQL tools and a knowledge management module into a unified framework.

a knowledge management module that analyzes historical data and extracts actionable insights to guide error correction, query rewriting, and semantic refinement. We leave the detailed introduction to the knowledge management module in Appendix A.1.

Specifically, we focus on three common issues encountered during SQL execution in real-world applications: 1) resolving execution failures caused by syntax errors; 2) rewriting inefficient queries that result in excessively long execution times; 3) modifying SQL to better align with user intent.

### 3 SPECIALIZED SQL TOOLS

This section introduces our core design principle—the fragment processing strategy—which enables fine-grained, modular analysis for complex SQL tasks. Built upon this foundation, we further present four specialized tools in SQLGOVERNOR, each targeting a key subtask. The following subsections will detail three of these tools, while the elaboration of the equivalence verifier is provided in Appendix A.2.

#### 3.1 FRAGMENT PROCESSING STRATEGY

Modern SQL queries, especially in OLAP workloads, often exhibit deeply nested structures with multiple layers of subqueries and Common Table Expressions (CTEs). To enable scalable and precise analysis, we propose a recursive fragment processing strategy that decomposes an input SQL query into smaller, self-contained fragments. Each fragment is analyzed independently under the same procedure, significantly reducing the reasoning complexity for LLM-based agents.

Concretely, the strategy operates as follows: the input query is first partitioned into a main query and a collection of CTEs. Each of these components is then recursively parsed to extract subqueries, which are likewise treated as fragments. For every fragment, the system performs a localized analysis and stores the result as a tuple of the form  $\langle \text{fragment\_id}, \text{analysis\_res} \rangle$ . The pseudocode in Algorithm 1 outlines the core idea for fragment processing.

#### 3.2 QUERY REWRITER

The Query Rewriter tool performs query optimization via a two-stage process: evaluation and rewriting. In the evaluation stage, the tool analyzes the input query to determine whether it exhibits inefficiencies. Specifically, it localizes the bottlenecks to specific fragments and proposes targeted rewriting strategies. In the rewriting stage, the tool applies the selected rewriting strategies to generate an optimized version of the original query. By leveraging the fragment-level structure produced during the fragment processing stage, the rewriting is both context-aware and fine-grained, ensuring that improvements are applied precisely without altering the query semantics.

**Evaluation** The evaluation process combines rule-based pattern matching with LLM-driven reasoning to deliver both precise and innovative optimization strategies—enhancing the overall efficiency of complex SQL queries.

162 If any rules match during the initial phase, their detailed descriptions are retrieved from the knowledge base and combined with the corresponding fragments into a structured prompt template (referred to as Scenario 1). If no rules are matched, the system evaluates the query against a set of “already efficient SQL” rules. A successful match indicates that the query is already optimal and does not require further rewriting efforts; otherwise, it suggests potential optimization opportunities that may not be covered by existing rules. In such cases, the query is passed to another predefined prompt template (referred to as Scenario 2).  
 163  
 164  
 165  
 166  
 167  
 168

169 Following the broad assessment, a detailed analysis is conducted using the LLM to refine and expand upon the rewriting suggestions. Specifically: (1) For Scenario 1, the prompt instructs the LLM to evaluate the applicability of each suggested rule, providing justifications and transforming general recommendations into actionable instructions where appropriate. (2) For Scenario 2, the prompt directs the LLM to analyze the intent of the original query and explore alternative, more efficient formulations that preserve semantic equivalence. The output of this stage is standardized in JSON format, enabling seamless integration with the subsequent rewriting module.  
 170  
 171  
 172  
 173  
 174  
 175

176 To clarify the “fragment processing” design, Listing 1 provides a representative example. This SQL  
 177 query contains six subqueries across different levels, with the deepest level of nesting being four.  
 178 We have numbered all subqueries according to the order in which they are analyzed. The analysis  
 179 results are as follows: Fragments 1-3 did not match any rules and meet the criteria for efficient SQL.  
 180 Fragment 4 matched the SAME\_TABLE\_JOIN rule, as it was detected to scan the same table (e.g.,  
 181 tb0) twice. Fragments 5-6 also did not match any rules and are considered efficient. The outermost  
 182 query (i.e., fragment 7) satisfies a rule involving a LEFT JOIN and an IS NOT NULL condition.  
 183 These two defects were further confirmed by the LLM.  
 184

185 Listing 1: Show the working mechanism of the fragment processing design.

```

1  -- Fragment 7, Line 2-27
2  SELECT tb0.c0,
3  -- Fragment 5, Line 4
4  (SELECT tb3.c1 - tb3.c2 FROM tb3 WHERE tb3.ds = tb1.ds),
5  -- Fragment 6, Line 6
6  (SELECT AVG(tb4.c3) FROM tb4 WHERE tb4.ds = tb1.ds AND tb4.c3 > 100)
7  FROM tb1
8  LEFT JOIN tb2 ON tb1.ds = tb2.ds
9  WHERE tb2.ds is NOT NULL AND
10 tb1.dcrs <=
11 (
12  -- Fragment 4, Line 13-27
13  SELECT IFNULL(t1.c1 / t2.c2, 100) AS dcrs
14  FROM
15  -- Fragment 2, Line 16-22
16  (SELECT MIN(c) AS c1
17  FROM
18  -- Fragment 1, Line 19-22
19  (SELECT COUNT(*) AS c, ds
20  FROM tb0
21  WHERE ds >= '1014' AND ds < '1016'
22  GROUP BY ds)) AS t1,
23  -- Fragment 3, Line 24-26
24  (SELECT COUNT(*) AS c2
25  FROM tb0
26  WHERE ds = '1016') AS t2
27 )

```

205  
 206 **Rewriting** As previously described, the Query Rewriter identifies potential inefficiencies in the  
 207 input SQL and generates a set of actionable rewriting suggestions during the evaluation phase. In  
 208 addition, the system retrieves relevant historical rewriting examples from the knowledge base by  
 209 matching the current SQL fragments and their associated optimization rules. The LLM then inte-  
 210 grates this information and synthesizes into a semantically equivalent yet execution-efficient SQL  
 211 query that incorporates the suggested optimizations.

212 Listing 2 illustrates the rewritten SQL for the query presented in Section 5.2.1. Specifically, the  
 213 pattern involving a LEFT JOIN combined with the IS NOT NULL condition is replaced with an  
 214 INNER JOIN. This transformation is effective because an INNER JOIN naturally filters out rows  
 215 with null values, thus achieving the same result as the original query but with a more efficient join  
 operation. Furthermore, the two separate scans of table tb0 are merged into a single scan to reduce

216 I/O overhead. The WHERE and SELECT clauses are appropriately adjusted to preserve the query’s  
 217 semantic correctness.  
 218

219 Listing 2: Show the rewritten SQL for the example in Listing 1.

```

220 1 WITH cte AS
221 2   (SELECT
222 3     IFNULL(MIN(CASE WHEN ds >= '1014' AND ds < '1016' THEN cnt END) /
223 4           SUM(CASE WHEN ds = '1016' THEN 1 ELSE 0 END), 100) AS dcrs
224 5   FROM (
225 6     SELECT ds, COUNT(*) AS cnt
226 7     FROM tb0
227 8     WHERE ds >= '1014' AND ds <= '1016'
228 9     GROUP BY ds
22910   )
22911   SELECT tb0.c0,
22912     (SELECT tb3.c1 - tb3.c2 FROM tb3 WHERE tb3.ds = tb1.ds),
22913     (SELECT AVG(tb4.c3) FROM tb4 WHERE tb4.ds = tb1.ds AND tb4.c3 > 100)
22914   FROM tb1
22915   INNER JOIN tb2 ON tb1.ds = tb2.ds
22916   WHERE tb1.dcrs <= (SELECT dcrs FROM cte)
230
231
232
233 3.3 QUERY MODIFIER
234
235 We classify modification requests into four general categories: (1) Realizing a specified semantics:  

  236 Adjust the SQL logic to align with the user’s intended meaning. (2) Explaining the SQL: Preserve  

  237 the original logic while adding comments or annotations for clarity. (3) Adopting a specified syntax:  

  238 Maintain semantic equivalence while adapting the query to match the user’s stylistic or structural  

  239 preferences. (4) Other SQL-related tasks: Capture queries that do not clearly fall into the above three  

  240 categories, such as stylish polishing. To fulfill each request, the tool follows a three-step pipeline:  

  241 (1) metadata preparation, (2) user intent clarification, and (3) query modification. Note that the  

  242 definitions of categories are rather flexible and can be customized.
243
244 Metadata Preparation During the metadata preparation stage, we extract the target SQL snippet  

  245 along with its surrounding context to provide a comprehensive view of the query environment. We  

  246 gather metadata from two primary sources. The first includes tables and columns referenced in  

  247 the SQL snippet, along with their names and descriptions. The second source is derived from the  

  248 user’s historical query logs, where we identify the top- $k$  most frequently accessed tables. For each  

  249 of these tables, we extract relevant metadata—such as schema information and usage patterns—to  

  250 help the LLM better understand the context and semantics of the query. Additionally, we append a  

  251 current timestamp to the metadata to provide temporal grounding, which is particularly useful when  

  252 handling evolving schema or time-sensitive queries.
253
254 User Intent Clarification The user intent clarification step maps the natural language request to  

  255 one of the four predefined categories described earlier. This classification combines two comple-  

  256 mentary strategies: heuristic keyword matching and semantic similarity scoring. In the heuristic  

  257 keyword matching strategy, we identify a set of domain-specific keywords and phrases that are com-  

  258 monly associated with each modification type. For each category  $C_j$ , we define a corresponding  

  259 keyword set  $\mathcal{KW}_j = \{k_{j1}, k_{j2}, \dots, k_{jn_j}\}$ . Given a user request  $\mathcal{Q}$ , we compute a weighted match-  

  260 ing score  $S_j^{kw}$  for each category as  $S_j^{kw} = \frac{1}{N_j} \sum_{k_{ji} \in \mathcal{KW}_j} \text{match}(\mathcal{Q}, k_{ji}) \times w_{ji}$ , where  $N_j$  is the  

  261 total number of candidate keywords in  $\mathcal{KW}_j$ ;  $\text{match}(\mathcal{Q}, k_{ji})$  is a binary function returning 1 if the  

  262 keyword  $k_{ji}$  appears in the request  $\mathcal{Q}$ , and 0 otherwise;  $w_{ji}$  denotes the weight assigned to keyword  

  263  $k_{ji}$ , reflecting its relative importance within category  $C_j$ .
264
265 To complement the keyword-based method, we also employ semantic embeddings to capture  

  266 more nuanced intent signals. We have explored two distinct embedding pathways: (1) Sentence-  

  267 Transformer with Masking: We pre-process the query by replacing specific metadata (e.g., table/-  

  268 column names) and constant values with special tokens [MASK], then encode the masked text  

  269 into a vector  $\mathbf{e}_{\mathcal{Q}}$  using a Sentence-Transformer model (e.g., RoBERTa Liu et al. (2019)). (2)  

  270 Instruction-aware Qwen3-Embedding: We utilize the Qwen3 embedding model Zhang et al. (2025)  

  271 to encode the original query  $\mathcal{Q}$ , guided by an instruction prompt that directs the model to focus  

  272 on the user’s action intent rather than details such as schema identifiers. During our development  

  273 and testing, the Instruction-Aware Qwen3 Embedding consistently outperforms the masking-based
274
```

270 Sentence-Transformer approach in both classification accuracy and robustness to domain variations.  
 271 This is attributed to its ability to better align with the LLM’s internal reasoning process and its use  
 272 of instruction-tuned representations that emphasize action-oriented semantics. What’s more, real-  
 273 world user queries often suffer from ambiguity, incomplete descriptions, or informal phrasing. In  
 274 such cases, rule-based detection and masking strategies tend to fail.

275 Once the embedding method has been selected, we construct a representative embedding vector  $\mathbf{e}_{C_j}$   
 276 for each category  $C_j$ . These are derived by collecting historical user queries, classifying their intents  
 277 using an LLM, and computing the centroid embedding for each category.

279 We employ cosine similarity between the query embedding  $\mathbf{e}_Q$  and the category centroid embed-  
 280 dings  $\mathbf{e}_{C_j}$  as a measure of semantic proximity. This semantic score is combined with the heuristic  
 281 keyword matching score to form the final classification decision. Formally, we define the final clas-  
 282 sification score  $F_j$  for each category  $C_j$  as  $F_j = \alpha \cdot S_j^{kw} + \beta \cdot \text{similarity}(\mathbf{e}_Q, \mathbf{e}_{C_j})$ , where  $\alpha$  and  $\beta$  are  
 283 weighting parameters used to balance the heuristic keyword matching score and the semantic simi-  
 284 larity score. To ensure robust and reliable intent clarification, we introduce a confidence threshold  
 285  $\theta$ . If the maximum classification score  $\max(F_j)$  across all categories falls below this threshold, the  
 286 system treats the request as unsupported and rejects the modification task. This mechanism helps  
 287 filter out ambiguous or outlier queries that do not align well with any known modification type,  
 288 thereby maintaining the integrity and reliability of the classification pipeline.

289 **Modification** Once the necessary data has been collected and the user’s intent has been classified,  
 290 we construct a structured prompt tailored to the identified modification type. The prompt integrates  
 291 the following components: the original SQL fragment, its surrounding context, relevant metadata  
 292 (e.g., schema information and top accessed tables), the current timestamp for temporal grounding,  
 293 and the natural language instruction from the user. This contextualized prompt is then fed into the  
 294 LLM, which reasons over the input and generates a modified SQL fragment that accurately fulfills  
 295 the user’s intent while preserving correctness and consistency.

### 297 3.4 SYNTAX ERROR CORRECTOR

299 The overall syntax correction workflow comprises three key stages: clarification, data preparation,  
 300 and correction.

302 **Clarification** The clarification stage begins by extracting structured information—such as the ex-  
 303 ception type, error location, and descriptive message—from DBMS error logs. This data is used for  
 304 embedding-based retrieval against a knowledge base of known error patterns and correction strate-  
 305 gies. Each retrieved strategy provides three key components to guide the LLM: (1) Schema De-  
 306 pendency: Indicates whether the error requires access to schema metadata (e.g., for column-related  
 307 errors) to avoid including large, complex schemas unnecessarily. (2) Correction Scope: Classifies  
 308 the error as either localized (e.g., a missing comma) or global, helping to focus the LLM’s attention.  
 309 (3) Correction Hints: Provides explicit, actionable guidance for the LLM, such as explaining the  
 310 root cause of a Column count mismatch error. This approach of selective schema inclusion  
 311 and localized correction keeps the prompt concise. This not only reduces inference latency and cost  
 312 but also prevents irrelevant information from interfering with the model’s reasoning capabilities.

314 **Data Preparation** Guided by the outputs from the clarification stage, the data preparation step  
 315 determines what information should be included in the final prompt. It selectively extracts relevant  
 316 schema components, or isolates specific query fragments for localized correction, depending on  
 317 the retrieved strategy. If the clarification stage fails to find a confident match in the knowledge  
 318 base, a conservative fallback strategy is applied: the full schema is retained, and a global correction  
 319 approach is used.

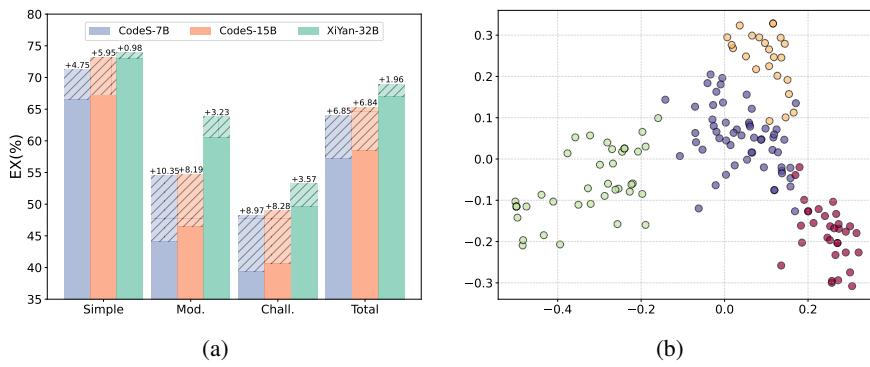
321 **Correction** In the final correction stage, the erroneous SQL (or fragment), along with its associ-  
 322 ated context—including selectively extracted schema information and tailored guidance—is assem-  
 323 bled into a structured prompt. The LLM agent then processes this input and generates a corrected  
 version of the SQL query.

## 324 4 EXPERIMENTS

326 This section presents the main results of our experiments on the BIRD and BIRD-CRITIC benchmarks,  
 327 along with a partial ablation study. For more comprehensive details on the datasets, experimental  
 328 settings, and additional results, please refer to the Appendix A.4- A.6.

### 330 4.1 MAIN RESULTS

332 **Results on BIRD** As SQLGOVERNOR is used as post-processing tool for the NL2SQL task, we  
 333 select three strong base models and one representative baseline. The base models are CodeS-7B,  
 334 CodeS-15B Li et al. (2024a) and XiYanSQL Gao et al. (2024b) and the baseline is *SQLFixAgent* Cen  
 335 et al. (2024).



348 Figure 2: (a) Performance on various difficulty data categories. (b) Visualization results of user  
 349 query clustering.

350 Table 4 presents a detailed comparison of SQLGOVERNOR’s performance against existing methods  
 351 on the BIRD dataset. When integrated with CodeS-15B, SQLGOVERNOR achieves EX and VES  
 352 scores of 65.32% and 67.87%, respectively, representing improvements of 6.84% in EX and 8.00%  
 353 in VES over the baseline. These gains outpace those of SQLFixAgent, which improves the base-  
 354 line by only 3% in EX and 4.35% in VES, highlighting SQLGOVERNOR’s superior effectiveness.  
 355 When paired with XiYan-32B, SQLGOVERNOR further enhances the already high baseline scores  
 356 of 68.97% (EX) and 70.89% (VES), achieving marginal but meaningful improvements of 1.96%  
 357 in EX and 3.10% in VES.

358 Furthermore, we carefully analyze the detailed performance of SQLGOVERNOR across various base  
 359 models. Figure 2a presents the results of this method on the three difficulty levels (Simple, Mod.,  
 360 and Chall.) of the BIRD dev set. The solid bars represent the results of the base model, while  
 361 the dashed bars above indicate the gains achieved by SQLGOVERNOR. It is evident that SQLGO-  
 362 VERNOR achieves significantly higher gains on the Mod. and Chall. difficulty levels compared  
 363 to Simple. Notably, on the CodeS-15B and XiYan-32B models, the metric gains for Chall. even  
 364 surpass those for Mod., making it the highest-performing category among the three difficulty levels,  
 365 with respective gains of +8.28% and +3.57%. This clearly demonstrates the advantage of SQLGO-  
 366 VERNOR in handling long and complex SQL queries.

366 In summary, SQLGOVERNOR exhibits strong performance improvements across all tested base-  
 367 line models, outperforming alternative methods such as SQLFixAgent and demonstrating promising  
 368 value even when applied to state-of-the-art models like XiYan-32B.

370 Table 1: Evaluation of SQLGOVERNOR on BIRD-CRITIC-Flash. Metric: SR (%).

| 372 Method           | 373 Category            |  |                         |  |                                 |
|----------------------|-------------------------|--|-------------------------|--|---------------------------------|
|                      | 374 Query               | 375 Management                         | 376 Personalization     | 377 Efficiency                         | 378 Total                       |
| 379 <b>Qwen3-32B</b> | 380 /                   | 381 18.8 /                             | 382 34.0 /              | 383 28.1 /                             | 384 22.7 /                      |
| 385 + SQLGOVERNOR    | 386 21.9 $\uparrow$ 3.1 | 387 <b>54.0</b> $\uparrow$ <b>20.0</b> | 388 35.9 $\uparrow$ 7.8 | 389 36.4 $\uparrow$ <b>13.7</b>        | 390 36.0 $\uparrow$ <b>10.0</b> |
| 391 <b>Qwen2.5</b>   | 392 /                   | 393 20.3 /                             | 394 46.0 /              | 395 32.8 /                             | 396 31.8 /                      |
| 397 -72B-Instruct    | 398 + SQLGOVERNOR       | 399 <b>26.6</b> $\uparrow$ <b>6.3</b>  | 400 52.0 $\uparrow$ 6.0 | 401 <b>43.8</b> $\uparrow$ <b>11.0</b> | 402 <b>45.5</b> $\uparrow$ 13.6 |
|                      |                         |  |                         |  | 403 <b>40.5</b> $\uparrow$ 8.5  |

379 **Results on BIRD-CRITIC-Flash** To evaluate the effectiveness of SQLGOVERNOR in addressing  
 380 SQL issues arising from user-provided natural language queries, we conduct experiments on the

378 BIRD-CRITIC-Flash benchmark. Our approach dynamically routes each SQL issue to the most  
 379 suitable tool: the Query Rewriter for efficiency-related problems, the Syntax Error Corrector for  
 380 execution errors, and the Query Modifier for other semantic or stylistic adjustments.  
 381

382 We evaluate two widely used LLMs—Qwen3-32B and Qwen2.5-72B-Instruct—in both the original  
 383 configurations provided by the benchmark team and with our toolkit integrated. As shown in Table 1,  
 384 SQLGOVERNOR consistently improves performance across all categories and both base models.  
 385

386 For Qwen3-32B, integrating SQLGOVERNOR leads to substantial gains, particularly in the Man-  
 387 agement category (+20.0%), where the success rate increases from 34.0% to 54.0%. Significant  
 388 improvements are also observed in Efficiency (+13.7%) and Personalization (+7.8%), indicating  
 389 that both query rewriting and semantic alignment benefit greatly from our toolkit. Overall, the total  
 390 score rises from 26.0% to 36.0%. When applied to Qwen2.5-72B-Instruct, SQLGOVERNOR still  
 391 delivers consistent gains. The largest improvement is seen in Personalization (+11.0%).  
 392

393 In addition to success rate, we measure the average end-to-end inference time per query on both  
 394 models. For Qwen3-32B, the runtime increases from 8.5s (base) to 18.4s with our toolkit; similarly,  
 395 for Qwen2.5-72B-Instruct, it increases from 9.8s (base) to 21.3s. While this represents a non-trivial  
 396 overhead, it is largely due to the multi-stage processing pipeline—including intricate problem anal-  
 397 ysis and solving process powered by LLM—that is essential for achieving high-quality corrections in  
 398 complex OLAP queries.  
 399

## 400 4.2 ABLATION STUDY

401 **User Intent Clarification** To evaluate the effectiveness of user intent clarification in the SQL  
 402 Modifier, we sampled 150 real-world query tasks from our production environment and manually  
 403 annotated them with intent categories. We then tested the performance of the Instruction-Aware  
 404 Qwen3 Embedding in classifying these intents. Specifically, we used an embedding model with 8B  
 405 parameters and a vector dimension of 1024. As a baseline, we also evaluated Qwen3-32B, where  
 406 the LLM directly performs classification without prior embedding-based filtering. Both models were  
 407 deployed under identical execution environments to ensure fair comparison.  
 408

409 The results are as follows: the Instruction-Aware Qwen3 Embedding achieves an accuracy of 78.9%,  
 410 with an average inference latency of 0.173 seconds. In contrast, Qwen3-32B achieves higher accu-  
 411 racy at 84.3%, but incurs a significantly higher average latency of 0.354 seconds. These findings  
 412 suggest that while the LLM-based classifier offers relatively higher accuracy (5.4%), the embedding-  
 413 based approach provides a favorable trade-off between speed and performance—making it particu-  
 414 larly suitable for high-throughput or latency-sensitive applications.  
 415

416 To provide a more intuitive understanding of the embedding quality, we applied PCA to reduce the  
 417 embedding vectors to two dimensions and visualized them using scatter plots, as shown in Figure 2b.  
 418

419 **Rewriting** To evaluate the performance of SQL rewriting tool, we used Payment-SQL, a test set  
 420 that closely aligns with OLAP scenario, and employed ETS and ETOG as evaluation metrics. We  
 421 selected four representative models as baselines: Qwen2.5-72B-Instruct Hui et al. (2024), Qwen3-  
 422 32B Zhang et al. (2025), LLM-R<sup>2</sup> Li et al. (2024c), and GenRewrite Liu & Mozafari (2024).  
 423 Specifically, Qwen2.5-72B-Instruct and Qwen3-32B are general-purpose LLMs that are instructed  
 424 to rewrite the input SQL query in a single inference step. LLM-R<sup>2</sup> employs LLMs to select ap-  
 425 propriate rewriting rules and trains a separate demonstration recommendation model to guide the  
 426 rewriting process. GenRewrite represents the first non-rule-based, end-to-end query rewriting ap-  
 427 proach that fully leverages the generative capabilities of LLMs. To ensure fair comparisons, we  
 428 maintained identical execution environments for the SQL queries before and after rewriting during  
 429 testing.  
 430

431 Table 2: Execution efficiency of SQLs in Payment-SQL after rewriting and time-cost for rewriting.  
 432

| 433 <b>Methods</b>     | 434 <b>ETOG(%)</b> | 435 <b>ETS(s)</b> | 436 <b>Cost(s)</b> | 437 <b><math>\Delta(s)</math></b>      |
|------------------------|--------------------|-------------------|--------------------|--|
| 438 Qwen3              | 439 11.06          | 440 20.19         | 441 15.24          | 442 $\uparrow 4.95$                    |
| 443 Qwen2.5            | 444 14.56          | 445 26.59         | 446 18.41          | 447 $\uparrow 8.18$                    |
| 448 LLM-R <sup>2</sup> | 449 29.87          | 450 54.53         | 451 46.85          | 452 $\uparrow 7.68$                    |
| 455 GenRewrite         | 456 31.25          | 457 57.07         | 458 38.36          | 459 $\uparrow 18.71$                   |
| 462 <b>SQLGOVERNOR</b> | <b>463 45.92</b>   | <b>464 83.86</b>  | <b>465 30.73</b>   | <b>466 <math>\uparrow 53.13</math></b> |

432 The results are presented in Table 2. From the table, we observe that SQLGOVERNOR exhibits a  
 433 significant performance advantage when applied to industrial-level OLAP workloads. On average,  
 434 across the entire test set, SQLGOVERNOR achieves a 45.92% reduction in execution time and an  
 435 83.86-second reduction in absolute execution time. We also report the rewriting cost, i.e., the time  
 436 required to perform the rewriting itself, and the net benefit ( $\Delta$ ), defined as the difference between  
 437 ETS and Cost. Notably, while SQLGOVERNOR incurs a relatively moderate rewriting overhead  
 438 (30.73s), it delivers the largest net benefit (+53.13s), demonstrating its practical viability in real-  
 439 world applications where query latency is critical.

440 Listing 3 presents an example of an rewritten SQL query from Payment-SQL. During the evaluation  
441 stage, the LLM provided the following rewriting suggestions: (1) Use the `WITH` clause to explicitly  
442 define the result of `UNION ALL` as a temporary table, making it easier for the rewriter to understand  
443 and optimize the query. (2) In the `UNION ALL` step, explicitly select only the columns that are  
444 actually needed, avoiding the retrieval and processing of unnecessary data.

Listing 3: Example of rewriting result from Payment-SQL.

```
447 -- Original SQL
448 SELECT AVG(duration)
449 FROM (
450     SELECT *, row_number() OVER (PARTITION by instanceid ORDER BY modifytime
451     DESC) AS id
452     FROM (
453         SELECT *
454         FROM table0
455         WHERE ds > '0201'
456         UNION ALL
457         SELECT *
458         FROM table1
459         WHERE ds > '0201'
460     ) a
461 ) b
462 WHERE id = 1 AND taskid IN (1, 12, 123) AND scriptid = 666
463 -- Rewritten SQL
464 WITH combined_data AS (
465     SELECT taskid, instanceid, scriptid, modifytime
466     FROM table0
467     WHERE ds > '0201'
468     UNION ALL
469     SELECT taskid, instanceid, scriptid, modifytime
470     FROM table1
471     WHERE ds > '0201'
472 )
473 SELECT AVG(duration)
474 FROM
475 (
476     SELECT *, ROW_NUMBER() OVER (PARTITION BY instanceid ORDER BY modifytime
477     DESC) AS id
478     FROM combined_data
479 ) b
480 WHERE id = 1 AND taskid IN (1, 12, 123) AND scriptid = 666
```

## 5 CONCLUSIONS

In this work, we present SQLGOVERNOR, the first comprehensive LLM-based SQL toolkit with integrated knowledge management. It unifies four core functionalities—syntax correction, query rewriting, semantic refinement, and consistency verification—into a single framework powered by a hybrid self-learning mechanism.

One of the key innovations lies in its fragment-wise processing strategy. By focusing on individual fragments such as subqueries and CTEs, the approach improves precision while reducing the cognitive burden on LLMs. Moreover, SQLGOVERNOR incorporates an expert-guided hybrid self-learning framework that continuously enhances performance by extracting patterns from execution outputs and validating generated rules with minimal expert input.

Extensive experiments show that SQLGOVERNOR consistently boosts base models' performance by up to 10% in key metrics on benchmarks like BIRD and BIRD-CRITIC. Deployed in production environments, it demonstrates strong utility and adaptability across real-world databases.

486 REFERENCES  
487

488 Qiushi Bai. *Improving SQL Performance Using Middleware-Based Query Rewriting*. PhD thesis,  
489 University of California, Irvine, 2023.

490 Qiushi Bai, Sadeem Alsudais, and Chen Li. Querybooster: Improving sql performance using mid-  
491 dleware services for human-centered query rewriting. *arXiv preprint arXiv:2305.08272*, 2023.

492 Mahdi Banisharif, Arman Mazloumzadeh, Mohammadreza Sharbaf, and Bahman Zamani. Auto-  
493 matic generation of business intelligence chatbot for organizations. In *2022 27th International  
494 Computer Conference, Computer Society of Iran (CSICC)*, 2022.

495 Michelangelo Ceci, Alfredo Cuzzocrea, and Donato Malerba. Effectively and efficiently supporting  
496 roll-up and drill-down olap operations over continuous dimensions via hierarchical clustering.  
497 *Journal of Intelligent Information Systems*, 44, 2015.

498 Jipeng Cen, Jiaxin Liu, Zhixu Li, and Jingjing Wang. Sqlfixagent: Towards semantic-accurate sql  
499 generation via multi-agent collaboration. *arXiv preprint arXiv:2406.13408*, 2024.

500 Hsinchun Chen, Roger HL Chiang, and Veda C Storey. Business intelligence and analytics: From  
501 big data to big impact. *MIS quarterly*, 2012.

502 Xinyun Chen, Maxwell Lin, Nathanael Schaefer, and Denny Zhou. Teaching large language models  
503 to self-debug. In *The 61st Annual Meeting Of The Association For Computational Linguistics*,  
504 2023.

505 Edgar F Codd. Providing olap (on-line analytical processing) to user-analysts: An it mandate.  
506 <http://www.arborsoft.com/papers/coddTOC.html>, 1993.

507 Haoran Ding, Zhaoguo Wang, Yicun Yang, Dexin Zhang, Zhenglin Xu, Haibo Chen, Ruzica Piskac,  
508 and Jinyang Li. Proving query equivalence using linear integer arithmetic. *Proceedings of the  
509 ACM on Management of Data*, 2023a.

510 Haoran Ding, Zhaoguo Wang, Yicun Yang, Dexin Zhang, Zhenglin Xu, Haibo Chen, Ruzica Piskac,  
511 and Jinyang Li. Proving query equivalence using linear integer arithmetic. *Proceedings of the  
512 ACM on Management of Data*, 2023b.

513 Rui Dong, Jie Liu, Yuxuan Zhu, Cong Yan, Barzan Mozafari, and Xinyu Wang. Slabcity: Whole-  
514 query optimization using program synthesis. *Proceedings of the VLDB Endowment*, 2023.

515 BV Elasticsearch. Elasticsearch. *software, version*, 2018.

516 Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou.  
517 Text-to-sql empowered by large language models: A benchmark evaluation. *Proceedings of the  
518 VLDB Endowment*, 2024a.

519 Yingqi Gao, Yifu Liu, Xiaoxia Li, Xiaorong Shi, Yin Zhu, Yiming Wang, Shiqi Li, Wei Li, Yuntao  
520 Hong, Zhiling Luo, Jinyang Gao, Liyu Mou, and Yu Li. Xiyan-sql: A multi-generator ensemble  
521 framework for text-to-sql. *arXiv preprint arXiv:2411.08599*, 2024b.

522 Yingqi Gao, Yifu Liu, Xiaoxia Li, Xiaorong Shi, Yin Zhu, Yiming Wang, Shiqi Li, Wei Li, Yuntao  
523 Hong, Zhiling Luo, et al. Xiyan-sql: A multi-generator ensemble framework for text-to-sql. *arXiv  
524 preprint arXiv:2411.08599*, 2024c.

525 Gartner. Gartner forecasts worldwide it spending to grow 3 percent in 2022, 2022. URL <https://www.gartner.com>.

526 Sneha Gathani, Peter Lim, and Leilani Battle. Debugging database queries: A survey of tools, tech-  
527 niques, and users. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing  
528 Systems*, 2020.

529 Goetz Graefe. The cascades framework for query optimization. *IEEE Data Eng. Bull.*, 1995.

530 Goetz Graefe and William J McKenna. The volcano optimizer generator: Extensibility and efficient  
531 search. In *Proceedings of IEEE 9th international conference on data engineering*, 1993.

540 Binyuan Hui, Jian Yang, Zeyu Cui, Jiaxi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang,  
 541 Bowen Yu, Kai Dang, et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*,  
 542 2024.

543 Andreas Kipf, Thomas Kipf, Bernhard Radke, Viktor Leis, Peter Boncz, and Alfons Kem-  
 544 per. Learned cardinalities: Estimating correlated joins with deep learning. *arXiv preprint*  
 545 *arXiv:1809.00677*, 2018.

546 Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo,  
 547 Hongcheng Gao, Wenjing Hu, Pengcheng Yin, Victor Zhong, Caiming Xiong, Ruoxi Sun, Qian  
 548 Liu, Sida Wang, and Tao Yu. Spider 2.0: Evaluating language models on real-world enterprise  
 549 text-to-sql workflows, 2024.

550 Alon Y Levy, Inderpal Singh Mumick, and Yehoshua Sagiv. Query optimization by predicate move-  
 551 around. In *VLDB*, 1994.

552 Haoyang Li, Jing Zhang, Hanbing Liu, Ju Fan, Xiaokang Zhang, Jun Zhu, Renjie Wei, Hongyan  
 553 Pan, Cuiping Li, and Hong Chen. Codes: Towards building open-source language models for  
 554 text-to-sql. *Proceedings of the ACM on Management of Data*, 2024a.

555 Jinyang Li, Binyuan Hui, Ge Qu, Jiaxi Yang, Binhu Li, Bowen Li, Bailin Wang, Bowen Qin,  
 556 Ruiying Geng, Nan Huo, et al. Can llm already serve as a database interface? a big bench for  
 557 large-scale database grounded text-to-sqls. *Advances in Neural Information Processing Systems*,  
 558 2024b.

559 Jinyang Li, Xiaolong Li, Ge Qu, Per Jacobsson, Bowen Qin, Binyuan Hui, Shuzheng Si, Nan Huo,  
 560 Xiaohan Xu, Yue Zhang, et al. Swe-sql: Illuminating llm pathways to solve user sql issues in  
 561 real-world applications. *arXiv preprint arXiv:2506.18951*, 2025.

562 Zhaodonghui Li, Haitao Yuan, Huiming Wang, Gao Cong, and Lidong Bing. Llm-r2: A large  
 563 language model enhanced rule-based rewrite system for boosting query efficiency. *arXiv preprint*  
 564 *arXiv:2404.12872*, 2024c.

565 Jie Liu and Barzan Mozafari. Query rewriting via large language models. *arXiv preprint*  
 566 *arXiv:2403.09060*, 2024.

567 Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike  
 568 Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining  
 569 approach. *arXiv preprint arXiv:1907.11692*, 2019.

570 Ryan Marcus, Parimarjan Negi, Hongzi Mao, Nesime Tatbul, Mohammad Alizadeh, and Tim  
 571 Kraska. Bao: Making learned query optimization practical. In *Proceedings of the 2021 Inter-  
 572 national Conference on Management of Data*, 2021.

573 Lili Mou, Ge Li, Lu Zhang, Tao Wang, and Zhi Jin. Convolutional neural networks over tree  
 574 structures for programming language processing. In *Proceedings of the AAAI conference on  
 575 artificial intelligence*, 2016.

576 M Muralikrishna et al. Improved unnesting algorithms for join aggregate sql queries. In *VLDB*,  
 577 1992.

578 Torben Bach Pedersen and Christian S Jensen. Multidimensional database technology. *Computer*,  
 579 2001.

580 Hamid Pirahesh, Joseph M Hellerstein, and Waqar Hasan. Extensible/rule based query rewrite  
 581 optimization in starburst. *ACM Sigmod Record*, 1992.

582 Mohammadreza Pourreza and Davood Rafiei. Din-sql: Decomposed in-context learning of text-to-  
 583 sql with self-correction. *Advances in Neural Information Processing Systems*, 2024.

584 Mohammadreza Pourreza, Hailong Li, Ruoxi Sun, Yeounoh Chung, Shayan Talaei, Gaurav Tarlok  
 585 Kakkar, Yu Gan, Amin Saberi, Fatma Ozcan, and Sercan O Arik. Chase-sql: Multi-path reasoning  
 586 and preference optimized candidate selection in text-to-sql. *arXiv preprint arXiv:2410.01943*,  
 587 2024.

594 Nadav Roiter. Boosting revenue and growth with a data flywheel, 2025.  
 595 URL [https://brightdata.com/blog/brightdata-in-practice/  
 596 using-data-flywheel-to-scale-your-business](https://brightdata.com/blog/brightdata-in-practice-using-data-flywheel-to-scale-your-business).

597 Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. Dbscan revisited,  
 598 revisited: why and how you should (still) use dbscan. *ACM Transactions on Database Systems  
 599 (TODS)*, 2017.

600 Ji Sun and Guoliang Li. An end-to-end learning-based cost estimator. *arXiv preprint  
 601 arXiv:1906.02560*, 2019.

603 Shayan Talaei, Mohammadreza Pourreza, Yu-Chen Chang, Azalia Mirhoseini, and Amin Saberi.  
 604 Chess: Contextual harnessing for efficient sql synthesis. *arXiv preprint arXiv:2405.16755*, 2024.

606 Stackoverflow team. Stack overflow developer survey 2023, 2023a. URL <https://survey.stackoverflow.co/2023/>.

608 StarRocks team. Starrocks, 2023b. URL <https://github.com/StarRocks/starrocks>.

610 Erik Thomsen. *OLAP solutions: building multidimensional information systems*. John Wiley &  
 611 Sons, 2002.

612 TPC. Tpc-h. URL <https://www.tpc.org/tpch/>.

614 Immanuel Trummer, Junxiong Wang, Ziyun Wei, Deepak Maram, Samuel Moseley, Saehan Jo,  
 615 Joseph Antonakakis, and Ankush Rayabhari. Skinnerdb: Regret-bounded query evaluation via  
 616 reinforcement learning. *ACM Transactions on Database Systems (TODS)*, 2021.

617 Panos Vassiliadis and Timos Sellis. A survey of logical models for olap databases. *ACM Sigmod  
 618 Record*, 1999.

620 Zhaoguo Wang, Zhou Zhou, Yicun Yang, Haoran Ding, Gansen Hu, Ding Ding, Chuzhe Tang, Haibo  
 621 Chen, and Jinyang Li. Wetune: Automatic discovery and verification of query rewrite rules. In  
 622 *Proceedings of the 2022 International Conference on Management of Data*, 2022.

623 Tianhao Wu, Weizhe Yuan, Olga Golovneva, Jing Xu, Yuandong Tian, Jiantao Jiao, Jason Weston,  
 624 and Sainbayar Sukhbaatar. Meta-rewarding language models: Self-improving alignment with  
 625 llm-as-a-meta-judge. *arXiv preprint arXiv:2407.19594*, 2024.

626 Zixuan Yi, Yao Tian, Zachary G Ives, and Ryan Marcus. Low rank learning for offline query  
 627 optimization. *arXiv preprint arXiv:2504.06399*, 2025.

629 Chaoqun Zhan, Maomeng Su, Chuangxian Wei, Xiaoqiang Peng, Liang Lin, Sheng Wang, Zhe  
 630 Chen, Feifei Li, Yue Pan, Fang Zheng, et al. Analyticdb: real-time olap database system at  
 631 alibaba cloud. *Proceedings of the VLDB Endowment*, 2019.

632 Yi Zhan, Longjie Cui, Han Weng, Guifeng Wang, Yu Tian, Boyi Liu, Yingxiang Yang, Xiaoming  
 633 Yin, Jiajun Xie, and Yang Sun. Towards database-free text-to-sql evaluation: A graph-based  
 634 metric for functional correctness. In *Proceedings of the 31st International Conference on Com-  
 635 putational Linguistics*, 2025.

637 Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie,  
 638 An Yang, Dayiheng Liu, Junyang Lin, et al. Qwen3 embedding: Advancing text embedding and  
 639 reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025.

640 Fuheng Zhao, Lawrence Lim, Ishtiyaque Ahmad, Divyakant Agrawal, and Amr El Abbadi. Llm-  
 641 sql-solver: Can llms determine sql equivalence? *arXiv preprint arXiv:2312.10321*, 2023.

642 Qi Zhou, Joy Arulraj, Shamkant Navathe, William Harris, and Dong Xu. Automated verification  
 643 of query equivalence using satisfiability modulo theories. *Proceedings of the VLDB Endowment*,  
 644 2019.

646 Qi Zhou, Joy Arulraj, Shamkant B Navathe, William Harris, and Jinpeng Wu. Spes: A symbolic  
 647 approach to proving query equivalence under bag semantics. In *2022 IEEE 38th International  
 Conference on Data Engineering (ICDE)*, 2022.

648 Xuanhe Zhou, Guoliang Li, Chengliang Chai, and Jianhua Feng. A learned query rewrite system  
 649 using monte carlo tree search. *Proceedings of the VLDB Endowment*, 2021.  
 650

651 Xuanhe Zhou, Zhaoyan Sun, and Guoliang Li. Db-gpt: Large language model meets database. *Data  
 652 Science and Engineering*, 2024.

653

## 654 A APPENDIX

### 655 A.1 KNOWLEDGE MANAGEMENT

658 To enhance the effectiveness of LLMs in SQL-related tasks, we design a knowledge management  
 659 module paired with dedicated recall and retrieving techniques. SQLGOVERNOR benefits from past  
 660 experiences and domain-specific insights to reduce redundant computations and improve both accu-  
 661 racy and consistency. To further automate knowledge acquisition and minimize reliance on manual  
 662 expert input, we implement structured knowledge bases inspired by the self-reinforcing data fly-  
 663 wheel mechanism Roiter (2025). Each knowledge base continuously accumulates and updates rules  
 664 from both successful and failure cases, enabling the system to improve over time with minimal  
 665 supervision.

666

| 667 Index                           | 668 Description   |
|-------------------------------------|---|
| 669 COUNT(DISTINCT)_SAME_FIELD      | 670 Pre-aggregate using a CTE and convert complex condition into 0/1 flags with the `MAX`<br>671 function and then use `SUM` to count these flags.  |
| 672 COUNT(DISTINCT)_DIFFERENT_FIELD | 673 Decompose multiple `COUNT(DISTINCT)` expressions within the same `SELECT` clause into<br>674 separate subqueries, each computing a single `COUNT(DISTINCT)` to enable parallel<br>675 execution. Subsequently, the results are combined using `JOIN` or `UNION ALL` operations. |
| 676 IN(SELECT)                      | 677 Rewrite the `IN (SELECT xxx)` clause as a `JOIN (SELECT xxx)` operation and apply the<br>678 `MAPJOIN` hint to fully load the subquery's corresponding small table into memory.   |

(a) “Rules”

| 679 Index                               | 680 Details   | 681 Tag                        |
|---|---|--------------------------------|
| 682 Embedding of the raw SQL's template | 683 Raw SQL:<br>684 <pre>SELECT * FROM table1 WHERE id IN (SELECT id FROM table2 WHERE cond = 0)</pre><br>685 Rewritten SQL:<br>686 <pre>SELECT /*+ MAPJOIN(t2) */ t1.* FROM table1 t1 JOIN (SELECT DISTINCT id FROM table2 WHERE cond = 0) t2 ON t1.id = t2.id</pre>   | 687 IN(SELECT)                 |
| 688 Embedding of the raw SQL's template | 689 Raw SQL:<br>690 <pre>SELECT ds, COUNT(DISTINCT filed_1) as unique_filed_1, COUNT(DISTINCT filed_2) as unique_filed_2, COUNT(DISTINCT filed_3) as unique_filed_3 FROM table GROUP BY ds</pre><br>691 Rewritten SQL:<br>692 <pre>WITH filed_1_count AS (SELECT ds, COUNT(DISTINCT filed_1) as unique_filed_1 FROM table GROUP BY ds), filed_2_count AS (SELECT ds, COUNT(DISTINCT filed_2) as unique_filed_2 FROM table GROUP BY ds), filed_3_count AS (SELECT ds, COUNT(DISTINCT filed_3) as unique_filed_3 FROM table GROUP BY ds) SELECT u.ds, u.unique_filed_1, o.unique_filed_2, p.unique_filed_3 FROM filed_1_count u JOIN filed_2_count o ON u.ds = o.ds JOIN filed_3_count p ON u.ds = p.ds</pre> | 693 COUNT(DISTINCT)_SAME_FIELD |

(b) “Historical Data”

690 Figure 3: Demonstration of the knowledge base serving the Query Rewriter tool.  
 691

692 **Structure of the Knowledge Base** The knowledge base is carefully structured to meet the diverse  
 693 needs of different SQL tools, with each tool having its own dedicated repository. Each repository  
 694 is divided into two sub-modules: “Rules” and “Historical Data”. The “Rules” sub-module contains  
 695 structured entries that provide actionable guidance for the LLM in addressing specific SQL-related  
 696 tasks. These rules include transformation patterns, rewriting strategies, and syntactic corrections,  
 697 enabling the model to apply well-defined solutions in a consistent manner. The “Historical Data”  
 698 sub-module stores high-quality, representative cases collected from real-world applications of the  
 699 SQL tools. By analyzing these past examples, the LLM can recognize recurring patterns and adopt  
 700 strategies that have proven effective in similar contexts.

701 Figure 3 illustrates the structure of the knowledge base used by the Query Rewriter. Each rule entry  
 702 consists of two fields: *index*, used for efficient retrieval, and *description*, which provides a detailed

702 explanation of the rule’s application and scope. Each historical data entry includes three fields: *index*  
 703 (for retrieval), *details* (a full description of the case), and *tag*, which links the case to relevant rule  
 704 indices, thereby facilitating cross-referencing between observed problems and applicable solutions.  
 705

706 **Knowledge Storage and Retrieval** To effectively support the diverse types of knowledge entries,  
 707 we design tailored storage and retrieval strategies for each sub-module of the knowledge base.  
 708 For the “Historical Data” module, all entries are stored in a vector database—specifically Star-  
 709 Rocks team (2023b). Retrieval is based on the structural and semantic similarity between SQL  
 710 templates. To facilitate this, we first extract the template of each SQL query by replacing concrete  
 711 identifiers (e.g., table and column names) with symbolic placeholders. This abstraction allows the  
 712 encoder to focus on high-level patterns rather than surface-level variations. During retrieval, the  
 713 input SQL is similarly transformed into a template, encoded into a vector representation, and used  
 714 to retrieve the top- $k$  most similar historical cases via cosine similarity. These candidates are further  
 715 filtered using the *Tag* field to ensure alignment with the specific task or error type.

716 In contrast, the “Rules” module employs different strategies depending on its application con-  
 717 text. For the Query Rewriter, each rule is associated with a unique label (e.g., `IN(SELECT)`)  
 718 that captures its applicability condition. We store these rules in an ElasticSearch Elastic-  
 719 search (2018) database to enable efficient exact matching during retrieval. For the Syntax Er-  
 720 rror Corrector, exception categories (e.g., `RuntimeException`, `SqlValidatorException`)  
 721 are often too coarse-grained to be informative. Therefore, we retain more detailed error  
 722 messages (e.g., `SqlValidatorException: INNER, LEFT, RIGHT or FULL join` requires a condition  
 723 (`NATURAL` keyword or `ON` or `USING` clause)) as re-  
 724 trieval targets. These messages are stored in a vector database. Given the verbosity of real-world  
 725 SQL execution logs, we first apply regular expressions to extract key information before encoding it  
 726 into vector representations for retrieval.

727 **Hybrid Self-Learning Mechanism** To ensure the knowledge base is both effective and actionable,  
 728 we adopt a multi-source initialization strategy tailored to each sub-module. For the “Rules” sub-  
 729 module, initialization involves aggregating knowledge from diverse sources and organizing it into  
 730 structured entries. For the Query Rewriter, the rule set is primarily sourced from domain experts and  
 731 is designed to complement the built-in optimization rules of the DBMS. These expert-defined rules  
 732 have been rigorously validated to ensure their practical effectiveness in real-world OLAP scenarios.  
 733 For the Syntax Error Corrector, the rule base is initialized using frequently asked questions (FAQs)  
 734 and common error-handling guidelines from technical documentation. These are formalized into  
 735 `exception, fixing action` pairs and further mapped to the `index, description` structure for retrieval  
 736 compatibility. For the “Historical Data” sub-module, initialization is conducted in two ways: (1)  
 737 manually curating high-quality cases that align with existing rules, and (2) extracting representative  
 738 queries from execution logs. These cases are selected based on criteria such as execution frequency,  
 739 complexity, and historical performance, ensuring their relevance and utility in future inference tasks.

740 Furthermore, both sub-modules support incremental updates through the self-learning mechanism,  
 741 allowing the system to refine its knowledge over time based on new data and user feedback. Specifi-  
 742 cally, we design a rule update paradigm that automatically extracts supplementary rules by analyzing  
 743 SQL queries that fail to meet user requirements—such as those resulting in execution errors, return-  
 744 ing incorrect results, or exhibiting excessive runtime. A heuristic pattern recognition-based data  
 745 filtering algorithm is first applied to extract relevant features (e.g., error messages, query structures,  
 746 execution times) from execution logs. An LLM agent then analyzes this information, identifies  
 747 common inefficiencies or mistakes, and generates structured knowledge entries. The prompt used  
 748 for rule generation is presented in Listing 4.

749 **Listing 4: Prompts for generating rules.**

750 **Task Description:**

751 You are provided with an SQL query along with its execution outputs (e.g. logs and  
 752 results) from DBMS. Your task is to analyze the logs and results to identify  
 753 potential errors or inefficiencies in the query.  
 754

755 **Instruction:**

- 756 1. Review the execution logs and results to determine whether the query contains errors  
 757 or inefficiencies.
- 758 2. For each confirmed problem, try to distill it into a generalizable rule, including the  
 759 abstract problem pattern, detailed description, and its solution.

756        3. Convert your findings into JSON format, where the key is the problem pattern used as  
 757        index and the value is a detailed description of the problem along with possible  
 758        corrective actions.  
 759        **Demonstration:**  
 760        {Few-shot examples curated from previously validated rule entries.}  
 761        **Question:**  
 762        {SQL query and user query if necessary.}  
 763        **Execution Outputs:**  
 764        {From logs and results.}

765        To ensure accuracy and reliability, experts verify the newly generated rules based on predefined con-  
 766        ditions. We implement a threshold-triggered mechanism to maintain the relevance of the knowledge  
 767        base. Verification is triggered when either: (a) the number of new rules exceeds a threshold  $t_1$ ,  
 768        or (b) the time elapsed since the last update surpasses a threshold  $t_2$ , where  $t_1 = \lfloor \lambda \cdot \sqrt{N_{\text{current}}} \rfloor$   
 769        and  $t_2 = \beta \cdot \mathbb{E}[\Delta t_{\text{historical}}]$ , with  $\lambda = 2.5$  controlling capacity scaling and  $\beta = 1.3$  for tempo-  
 770        ral adaptation. Here  $N_{\text{current}}$  denotes the current number of rules, and  $\mathbb{E}[\Delta t_{\text{historical}}]$  represents the  
 771        expected historical update interval. To avoid redundancy, semantically equivalent rules are iden-  
 772        tified and clustered through the following process: (a) The description field of rule  $r_i$  is encoded  
 773        using RoBERTa-base Liu et al. (2019). (b) Compute the pairwise similarity. (c) Rules are grouped  
 774        around  $r_i$  into  $C_k(i)$  based on the similarity score using DBSCAN Schubert et al. (2017). (d) Merge  
 775        semantically equivalent rules using centroid synthesis.

$$\mathbf{e}_i = \text{RoBERTa}(r_i.\text{get(description)}) \in \mathbb{R}^{768} \quad (1)$$

$$s(r_i, r_j) = \frac{\mathbf{e}_i^\top \mathbf{e}_j}{\|\mathbf{e}_i\| \|\mathbf{e}_j\|} \quad (2)$$

$$r_{\text{new}} = \arg \min_{r \in C_k(i)} \sum_{r_i \in C_k(i)} \|\mathbf{e}_r - \mathbf{e}_{r_i}\|_2 \quad (3)$$

782        The updated rules are then applied to subsequent tasks, generating new data that perpetuates the  
 783        improvement cycle.

784        The update mechanism for the “Historical Data” sub-module is relatively straightforward. As men-  
 785        tioned earlier, for rules that have been validated by experts, we incorporate the corresponding in-  
 786        stances into the historical data. This continuous cycle of knowledge collection, validation, and  
 787        application ensures that the knowledge base remains up-to-date and effective, thereby enhancing the  
 788        overall performance of the LLM-based SQL toolkit.

## 789        A.2 EQUIVALENCE VERIFIER

791        We conduct a systematic review of representative SQL equivalence verification methods and sum-  
 792        marize their characteristics in Table 3.

794        Formal methods such as SPES Zhou et al. (2022) offer rigorous correctness guarantees through sym-  
 795        bolic execution but are limited in practical applicability due to type constraints and poor performance  
 796        on real-world benchmarks like TPC-H TPC. Graph-based approaches (e.g., FuncEvalGMN Zhan  
 797        et al. (2025)) achieve broader coverage via structural matching but require extensive model training,  
 798        leading to high deployment costs and limited generalization across diverse query patterns. LLM-  
 799        based solutions Zhao et al. (2023) employ advanced prompting techniques but still suffer from  
 800        limited accuracy on complex queries and exhibit positive bias in equivalence judgments Wu et al.  
 801        (2024).

802        Table 3: Comparison of representative SQL equivalence verification methods.

| 803 <b>Method</b>                            | 804 <b>Technical Basis</b> | 805 <b>Correctness Guarantee</b> | 806 <b>Applicable Scope</b> | 807 <b>Deployment Cost</b> |
|--|----------------------------|----------------------------------|-----------------------------|----------------------------|
| 808        SPES Zhou et al. (2022)           | Symbolic Execution         | ✓                                | Very Limited                | Low                        |
| 809        SQLSolver Ding et al. (2023a)     | Formal Logic               | ✓                                | Limited                     | Low                        |
| 810        FuncEvalGMN Zhan et al. (2025)    | Graph Matching             | ✗                                | General                     | High                       |
| 811        LLM-SQL-Solver Zhao et al. (2023) | Probabilistic LM           | ✗                                | General                     | Medium                     |
| 812        SQLGOVERNOR                       | Probabilistic LM           | ✗                                | SELECT-based DML            | Medium                     |

813        To address these limitations, we propose a structured framework for SQL semantic equivalence ver-  
 814        ification, specifically tailored for SELECT-based DML queries in OLAP scenarios. Our approach

810 decomposes the verification process into two stages: (1) semantic intent extraction, which captures the meaning of each query using controlled LLM-mediated interpretation, and (2) hierarchical consistency checking, which performs field-level alignment and derivation trace analysis to assess equivalence.

811 In the first stage, each SQL query is translated into a structured natural language representation that captures the provenance of every field in the main `SELECT` clause. This includes source tables, transformation logic (e.g., aggregations or arithmetic expressions), and filtering or joining conditions. To enhance interpretability, we adopt a stepwise parsing strategy that starts with the innermost subqueries and progressively builds up the outer query semantics.

812 A lightweight pre-filtering module eliminates clearly inconsistent query pairs based on schema-level heuristics, such as mismatches in field count or source tables. For remaining pairs, an LLM agent performs detailed consistency analysis by aligning corresponding `SELECT` fields and reasoning about semantic equivalence.

813 Our prompt design enforces bidirectional field mapping, supports counterexample generation for non-equivalent pairs, and incorporates calibrated confidence scoring—ensuring both the interpretability and reliability of the verification process.

814

### 815 A.3 RELATED WORK

816

**817 Query Rewriting** Query rewriting can be classified according to the stages of the SQL query lifecycle. A typical query goes through several phases: parsing for syntax validation, binding for schema resolution, optimization using a cost model, and execution by the database engine. Based on this lifecycle, rewriting can be applied at different levels: (1) raw SQL queries before parsing, (2) logical plans generated by the binder, or (3) physical plans produced by the optimizer within the DBMS.

818

819 Physical plan rewriting focuses on selecting the most efficient execution plan among functionally equivalent alternatives. The DBMS optimizer employs search algorithms—such as dynamic programming Graefe (1995); Graefe & McKenna (1993)—to explore the space of physical plans without exhaustive enumeration. A cost model is used to estimate the execution cost of each candidate plan. Recent approaches have introduced machine learning and deep learning techniques Kipf et al. (2018); Trummer et al. (2021); Sun & Li (2019) to enhance cost estimation and plan selection. However, these methods often suffer from long training times, limited adaptability across workloads, and high integration costs.

820

821 An alternative approach that preserves the existing optimizer architecture is Bao Marcus et al. (2021), which integrates a Tree Convolutional Neural Network Mou et al. (2016) with Thompson sampling to guide the selection of better hint sets. This hybrid approach improves plan generation without requiring a complete redesign of the optimizer, thereby balancing innovation with system compatibility.

822

823 Logical plan rewriting involves transforming a tree-structured representation of query operations, focusing on what to compute rather than how to compute it. This process is primarily driven by rule-based pattern matching, where heuristic rules—such as predicate push-down and operation merging—guide the transformation Levy et al. (1994); Pirahesh et al. (1992); Muralikrishna et al. (1992). Recent efforts aim to automate rule discovery. WeTune Wang et al. (2022) uses brute-force enumeration to identify and validate new rules, enhancing the internal optimizer’s capabilities. QueryBooster Bai et al. (2023) introduces a connector for user-defined rules, enabling task-specific rewriting strategies. Traditional rule application orders are often fixed and suboptimal. LR Zhou et al. (2021) applies Monte Carlo Tree Search to explore effective rewriting sequences, while LLM-R<sup>2</sup> Li et al. (2024c) leverages large language models to recommend context-aware rewriting rules, improving adaptability and generalization.

824

825 End-to-end SQL rewriting aims to enhance transparency and usability by rewriting queries before they enter the DBMS pipeline. This approach enables holistic transformations and avoids the limitations of local rewriting efforts. Recent studies have explored the use of Large Language Models (LLMs) to facilitate this process. The DB-GPT framework Zhou et al. (2024) categorizes such approaches into three paradigms: in-context learning, LLM fine-tuning, and DB-specific pre-training. GenRewrite Liu & Mozafari (2024), a representative in-context learning method, designs prompts to

864 guide LLMs in SQL rewriting and stores generated natural language rules in the NLR2s repository  
 865 for reuse. To mitigate LLM hallucinations, GenRewrite includes a validation and correction step to  
 866 ensure the reliability of the rewritten queries.

867 Additionally, middleware-based rewriting has been explored to offload optimization tasks from the  
 868 DBMS. This approach provides a flexible layer between the application and the database, enabling  
 869 query transformation before execution Bai (2023). Similarly, query rewriting has been integrated  
 870 into human-in-the-loop systems to support interactive data exploration, where users can iteratively  
 871 refine queries based on intermediate results.

873 **Query Error Detection and Correction** SQL query errors fall into two categories: syntax errors  
 874 and semantic errors. Syntax errors occur when a query violates SQL’s syntactical rules, preventing  
 875 execution. Traditional debugging methods, as noted by Gathani et al. Gathani et al. (2020), lack  
 876 automated correction and instead help users identify errors through techniques like visualizing in-  
 877 termediate results. Semantic errors arise when a query fails to return expected data, indicating a  
 878 mismatch between the query’s output and the user’s intent. Verifying query equivalence is crucial  
 879 in NL2SQL conversion Pourreza et al. (2024); Talaei et al. (2024); Gao et al. (2024c) and query  
 880 rewriting models Dong et al. (2023); Liu & Mozafari (2024); Wang et al. (2022). Existing SQL  
 881 equivalence provers use algebraic representations to verify query equivalence by solving mathemat-  
 882 ical problems Ding et al. (2023b); Zhou et al. (2019), offering high reliability but at high compu-  
 883 tational cost. For loosely bounded verification, heuristic rules and counterexample construction are  
 884 employed Dong et al. (2023), while some studies leverage LLMs for reasoning and judgment Liu &  
 885 Mozafari (2024).

---

886 **Algorithm 1:** Fragment Processing Strategy

---

888 **Input:** SQL query  $Q$   
 889 **Output:** Analysis result set  $S$

```

890 1  $S \leftarrow \emptyset$ ;                                // Initialize result set
891 2  $D \leftarrow \emptyset$ 
892 3  $Q_{main}, \{Q_{cte_j}\} \leftarrow \text{divide\_CTE}(Q)$ 
893 4 if  $Q_{main} = \emptyset$  then
894 5   return  $\emptyset$ 
895 6 end
896 7  $\{Q_{sub_i}\} \leftarrow \text{parse\_subqueries}(Q_{main})$ 
897 8 foreach  $Q_{sub} \in \{Q_{sub_i}\}$  do
898 9    $S_{sub} \leftarrow \text{fragment\_processing}(Q_{sub})$ 
900 10   $S \leftarrow S \cup S_{sub}$ 
901 11 end
902 12 foreach  $Q_{cte} \in \{Q_{cte_j}\}$  do
903 13    $S_{cte} \leftarrow \text{fragment\_processing}(Q_{cte})$ 
904 14    $S \leftarrow S \cup S_{cte}$ 
905 15 end
906   // Main query analysis (details omitted)
907 16 return  $S$ 

```

---

908 A.4 DATASETS AND METRICS

909 **Datasets.** To comprehensively evaluate the performance of SQLGOVERNOR, we select two rep-  
 910 resentative benchmarks: BIRD-CRITIC Li et al. (2025) and BIRD Li et al. (2024b). Additionally,  
 911 we have constructed a new dataset named Payment-SQL, which comprises analytical SQL queries  
 912 derived from real industrial scenarios, specifically designed to evaluate performance in handling  
 913 complex and diverse queries.

914 **BIRD-CRITIC** is an innovative SQL benchmark crafted to evaluate the critical capabilities of LLMs  
 915 in diagnosing and resolving user issues within real-world database environments. The benchmark  
 916 categorizes issues into four domains: *Query*, *Management*, *Personalization*, and *Efficiency*. These  
 917 categories align with the core functionalities of SQLGOVERNOR. For our experiments, we utilize

918 a light version, `bird-critic-1.0-flash-exp`, which consists of 200 user issues on Post-  
 919 greSQL.  
 920

921 **BIRD** serves as a challenging large-scale database text-to-SQL evaluation benchmark, designed to  
 922 bridge the gap between academic research and practical applications. It encompasses 95 extensive  
 923 databases and high-quality text-SQL pairs, with data storage reaching up to 33.4GB, spanning 37  
 924 professional fields. The validation set includes 1,534 test entries, offering a comprehensive evalua-  
 925 tion of text-to-SQL translation capabilities. Notably, in utilizing this dataset, we employ SQLGo-  
 926 VERNOR as a post-processing tool for NL2SQL models, aimed at further enhancing the quality of  
 927 generated SQL queries.  
 928

929 **Payment-SQL dataset** originates from real-world industrial OLAP scenarios and is curated by hu-  
 930 man experts based on execution logs. It contains 50 SQL queries, each involving an average of  
 931 2 tables and 11 columns, drawn from a schema of 74 tables with thousands of fields. Designed  
 932 specifically for evaluating SQL rewriting systems, Payment-SQL measures effectiveness through  
 933 execution time comparisons before and after rewriting in the same environment—directly reflect-  
 934 ing real-world performance gains. A key feature of Payment-SQL is its complexity: the average  
 935 query length is 421 tokens, far exceeding that of BIRD’s challenging category (107 tokens). Ac-  
 936 cording to Spider 2.0 Lei et al. (2024), where queries over 160 tokens are considered difficult, even  
 937 the shortest query in Payment-SQL (173 tokens) qualifies as hard, with the longest reaching 1169  
 938 tokens. This makes Payment-SQL a rigorous and realistic benchmark for evaluating the robustness  
 939 and scalability of SQL rewriting techniques in industrial applications. The dataset is available at  
 940 <https://anonymous.4open.science/r/SQLGovernor-33DF>.  
 941

942 **Evaluation Metrics.** On the BIRD-CRITIC-FLASH dataset, we follow the official guidelines and  
 943 use the success rate (SR) as the metric, as it effectively evaluates multiple aspects of performance  
 944 due to the well-designed test cases. For the BIRD dataset, we employ both Execution Accuracy  
 945 (EX) and Valid Efficiency Score (VES) metrics to comprehensively evaluate performance. In the  
 946 case of the Payment-SQL dataset, rewriting effectiveness is assessed using Execution Time Saved  
 947 (ETS) and Execution Time Optimization Gain (ETOG), calculated as follows:  
 948

$$949 \text{ETS} = ET_{\text{pre}} - ET_{\text{post}}, \text{ETOG} = \frac{\text{ETS}}{ET_{\text{pre}}} \times 100\%, \quad (4)$$

950 where  $ET_{\text{pre}}$  represents the execution time before rewriting and  $ET_{\text{post}}$  represents the execution  
 951 time after rewriting. It is worth noting that when using ETS and ETOG to evaluate SQL rewriting  
 952 tasks, we typically execute both the pre-optimized and post-optimized SQL queries in the same  
 953 system while excluding interference factors such as execution caching to ensure the objectivity and  
 954 reliability of the test results.  
 955

## 956 A.5 MORE ABLATION STUDY

957 **Error Correction** To evaluate the error correction capabilities of SQLGOVERNOR, we collect a  
 958 set of syntactically and semantically incorrect SQL queries generated by two strong LLM-based  
 959 NL2SQL systems—CodeS-7B and CodeS-15B—on the BIRD dataset. Queries that failed to exe-  
 960 cute due to syntax errors are fed into the Syntax Error Corrector, while those exhibiting semantic  
 961 misalignment are routed to the Query Modifier for refinement.  
 962

963 Table 5 presents the results of the error correction capabilities in SQLGOVERNOR. The findings  
 964 indicate that the module demonstrates strong error correction performance on the BIRD datasets,  
 965 as evidenced by the predictive results from both baseline models. For the CodeS-7B model, we  
 966 analyzed 691 erroneous cases, yielding an overall EX rate of 25.8%. Performance across difficulty  
 967 levels shows EX rates of 26.4% for simple cases, 26.2% for moderate cases, and 22.0% for challeng-  
 968 ing cases. In contrast, the CodeS-15B model, evaluated on 667 erroneous cases, achieved an overall  
 969 EX rate of 25.2%, with rates of 26.9% for simple cases, 23.1% for moderate cases, and 25.0% for  
 970 challenging cases.  
 971

972 **Equivalence Verification** We use the predictive results from the CodeS-7B model alongside  
 973 golden SQL queries to establish positive and negative pairs. Correctly predicted SQL queries are  
 974 classified as equivalent with the golden SQL (labeled as true), while incorrectly predicted queries  
 975 are deemed nonequivalent (labeled as false). The results are presented in Table 6. We report two  
 976

Table 4: Evaluation on BIRD’s dev set.

| Methods                                  | Dev set                          |                                  |
|--|----------------------------------|----------------------------------|
|  | EX(%)                            | VES(%)                           |
| <b>Prompt-based base models</b>          |                                  |                                  |
| Codex Li et al. (2024b)                  | 34.35                            | 43.41                            |
| ChatGPT Li et al. (2024b)                | 37.22                            | 43.81                            |
| GPT-4 Li et al. (2024b)                  | 46.35                            | 49.77                            |
| DIN-SQL + GPT-4 Pourreza & Rafiei (2024) | 50.72                            | 58.79                            |
| DAIL-SQL + GPT-4 Gao et al. (2024a)      | 54.76                            | 56.08                            |
| <b>Fine-tuning-based base models</b>     |                                  |                                  |
| T5-3B Li et al. (2024b)                  | 23.34                            | 25.57                            |
| CodeS-7B Li et al. (2024a)               | 57.17                            | 58.80                            |
| CodeS-15B Li et al. (2024a)              | 58.48                            | 59.87                            |
| XiYan-32B Gao et al. (2024b)             | 67.01                            | 67.79                            |
| <b>With post-processing tools</b>        |                                  |                                  |
| CodeS-7B + SQLFixAgent Cen et al. (2024) | 60.17 ( $\uparrow$ 3.00)         | 63.15 ( $\uparrow$ 4.35)         |
| CodeS-7B + SQLGOVERNOR                   | 64.02 ( $\uparrow$ <b>6.85</b> ) | 64.72 ( $\uparrow$ 5.92)         |
| CodeS-15B + SQLGOVERNOR                  | 65.32 ( $\uparrow$ 6.84)         | 67.87 ( $\uparrow$ <b>8.00</b> ) |
| XiYan-32B + SQLGOVERNOR                  | <b>68.97</b> ( $\uparrow$ 1.96)  | <b>70.89</b> ( $\uparrow$ 3.10)  |

Table 5: Error correction performance on the BIRD’s dev set including syntactic and semantic levels.

| Error Data Statistics        | Total | Simple | Mod. | Chall. |
|------------------------------|-------|--------|------|--------|
| <b>#CodeS-7B Error Case</b>  | 691   | 333    | 267  | 91     |
| <b>EX(%)</b>                 | 25.8  | 26.4   | 26.2 | 22.0   |
| <b>#CodeS-15B Error Case</b> | 667   | 324    | 255  | 88     |
| <b>EX(%)</b>                 | 25.2  | 26.9   | 23.1 | 25.0   |

metrics-accuracy and F1 score-and present the results in Table 6, the overall accuracy for verification is 78.9% and F1 score is 79.3%, indicating effective performance of the Verifier. It is noteworthy that the scores for challenging queries are lower than those for simpler queries, which is expected given the increased complexity of the SQL statements.

Table 6: Equivalence verification performance on the predictive results of CodeS-7B.

| Data Category             | Total | Simple | Mod. | Chall. |
|---------------------------|-------|--------|------|--------|
| <b>Verif. Accuracy(%)</b> | 78.9  | 81.2   | 76.9 | 71.0   |
| <b>Verif. F1(%)</b>       | 79.3  | 84.3   | 69.9 | 57.1   |

## A.6 DETAILED EXPERIMENTAL ANALYSIS

**Capability in processing long and complex SQL** We analyze the capability of SQLGOVERNOR in handling long SQL queries from two common scenarios: error correction and rewriting. Figure 2a presents the detailed performance of SQLGOVERNOR when using CodeS-7B, CodeS-15B, and XiYan-32B as base models. The shaded bars illustrate the performance improvement achieved by SQLGOVERNOR over the base models. SQLGOVERNOR consistently outperforms the base models across all categories, with particularly notable gains in the Challenge SQL section of the BIRD dataset.

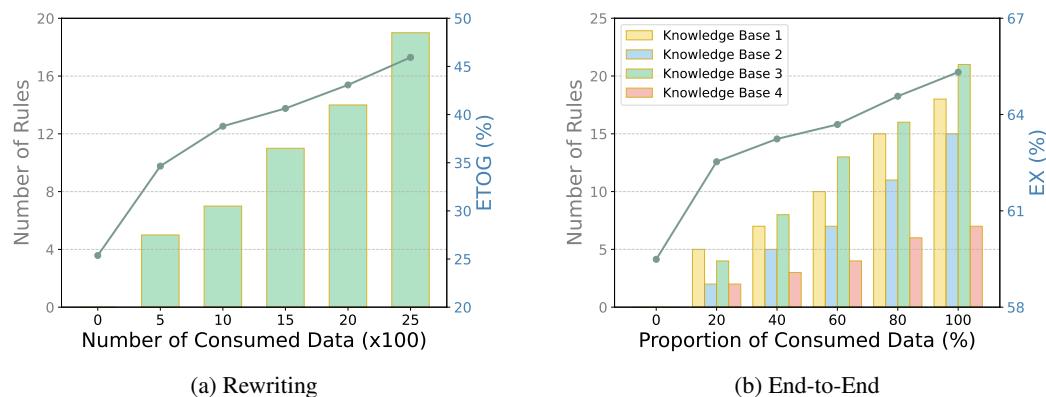
Furthermore, Table 2 illustrates the rewriting performance of SQLGOVERNOR on the industrial-level dataset Payment-SQL. Compared to general-purpose LLMs, SQLGOVERNOR exhibits a clear advantage in SQL rewriting tasks. Notably, the average token length of Payment-SQL reaches 421, far exceeding the complexity of SQL queries in the BIRD dataset. Additionally, all SQL queries in Payment-SQL meet the Hard SQL standard defined by Spider 2.0 Lei et al. (2024) (token length  $i$

1026 160). These results strongly demonstrate the superior rewriting capability of SQLGOVERNOR in  
 1027 handling long and complex SQL queries.  
 1028

1029 **Validation of productivity improvement** To validate the effectiveness of SQLGOVERNOR in ad-  
 1030 dressing productivity bottlenecks caused by fragmented SQL tool-chains, we conducted a controlled  
 1031 A/B testing with 60 practitioners from in-production data platform. Participants were stratified by  
 1032 SQL expertise (30 experts and 30 non-experts) and uniformly assigned to two groups-Group A: Ut-  
 1033 ilizing the integrated SQLGOVERNOR framework; Group B: Operating equivalent discrete modules  
 1034 through manual orchestration. Each subject executed 50 standardized SQL governance tasks span-  
 1035 ning evaluation, correction, rewriting, verification. We systematically measured-task completion  
 1036 time and tool-switching frequency. Results demonstrated statistically significant advantages for the  
 1037 integrated framework. Group A achieved 33% faster task completion, with non-experts exhibiting  
 1038 greater efficiency gains (41% improvement) compared to 25% for experts. This disparity corre-  
 1039 lates with Group B’s tool-switching patterns, where practitioners incurred 18% temporal overhead  
 1040 reconstructing workflow contexts between discrete modules. The empirical evidence quantitatively  
 1041 confirms that SQLGOVERNOR’s unified pipeline effectively mitigates fragmentation-induced pro-  
 1042 ductivity loss, particularly benefiting non-specialist users.  
 1043

1043 **“Evolving with every step”** To validate the effectiveness of our expert-guided hybrid self-learning  
 1044 mechanism in continuously enhancing the performance of SQLGOVERNOR across various SQL-  
 1045 related tasks, we collect and retain results at different stages for an end-to-end task and an individual  
 1046 tasks. This approach allows us to assess how SQLGOVERNOR improves its capabilities through  
 1047 self-learning.

1048 Specifically, for the end-to-end task, we use the predictive results of CodeS-15B on the BIRD’s dev  
 1049 set. For SQL rewriting task, we choose the Payment-SQL dataset to examine the iterative gains  
 1050 of SQLGOVERNOR in long SQL rewriting scenarios. The experimental results shown in Figure 4  
 1051 demonstrate that the hybrid self-learning approach not only enhances the performance of SQLGo-  
 1052 VERNOR but also provides a reliable foundation for its continuous rewriting in real-world industrial  
 1053 applications. Moreover, the effectiveness of this mechanism further validates the feasibility of tran-  
 1054 sitioning from expert-centric knowledge base construction to an expert-guided hybrid self-learning  
 1055 framework, thereby providing methodological support for reducing the cost of complete reliance on  
 1056 experts for knowledge collection and maintenance.



1069 Figure 4: The performance metrics of SQLGOVERNOR across different stages of self-learning. The  
 1070 bar chart corresponds to the left y-axis, while the line chart corresponds to the right y-axis.  
 1071

1072  
 1073  
 1074  
 1075  
 1076  
 1077  
 1078  
 1079