



# SQL-Factory: A Multi-Agent Framework for High-Quality and Large-Scale SQL Generation

Jiahui Li  
Zhejiang University  
li.jiahui@zju.edu.cn

Tongwang Wu  
Zhejiang University  
tongwang.wu@zju.edu.cn

Yuren Mao<sup>†</sup>  
Zhejiang University  
yuren.mao@zju.edu.cn

Yunjun Gao  
Zhejiang University  
gaoyj@zju.edu.cn

Yajie Feng  
Global Technical Service Dept,  
Huawei Technologies  
fengyajie@huawei.com

Huaizhong Liu  
Global Technical Service Dept,  
Huawei Technologies  
liuhuaizhong@huawei.com

## ABSTRACT

High quality SQL corpus is essential for intelligent databases. For example, Text-to-SQL requires SQL queries and corresponding natural language questions as training samples. However, collecting such a query corpus remains challenging in practice due to the high cost of manual annotation, which highlights the importance of automatic SQL generation. Despite recent advances, existing generation methods still face limitations in achieving both diversity and cost-effectiveness. Besides, many methods also treat all tables equally, which overlooks schema complexity and leads to under-utilization of structurally rich tables. To address these issues, this paper proposes a multi-agent framework for high-quality and large-scale SQL generation, dubbed SQL-Factory. It decomposes the generation process into three collaborative teams: the Generation Team explores diverse query structures using a powerful language model, the Expansion Team scales promising patterns via a lightweight language model, and the Management Team adaptively schedules the workflow and evaluates the quality of synthesized queries. This modular framework ensures a balanced trade-off between diversity, scalability, and generation cost. We apply SQL-Factory to four widely used benchmarks and generate over 300,000 SQL queries with less than \$200 API cost. Our generated queries achieve higher diversity compared to other methods, and extensive experiments demonstrate that the generated queries significantly improve model performance in various downstream tasks.

### PVLDB Reference Format:

Jiahui Li, Tongwang Wu, Yuren Mao, Yunjun Gao, Yajie Feng, and Huaizhong Liu. SQL-Factory: A Multi-Agent Framework for High-Quality and Large-Scale SQL Generation. PVLDB, 19(3): 292 - 305, 2025. doi:10.14778/3778092.3778093

### PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/LJHZju/SQL-Factory>.

<sup>†</sup>Yuren Mao is the corresponding author.

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](mailto: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. 3 ISSN 2150-8097. doi:10.14778/3778092.3778093




Figure 1: A conceptual illustration of SQL-Factory. The framework consists of six agents: (1) Table Selection Agent, (2) Generation Agent, (3) Management Agent, (4) Critical Agent, (5) Seed Selection Agent and (6) Expansion Agent.

## 1 INTRODUCTION

A high-quality SQL corpus is essential for advancing intelligent database technologies, such as cardinality estimation [31, 38, 53], configuration tuning [5, 30, 51, 63] and natural language interfaces [16, 28, 48]. For example, to train a learning-based configuration tuning system, it is necessary to observe how different SQL queries interact with various configuration settings, which requires a large number of diverse SQL workloads. Besides, since SQL can be easily translated into natural language, it also facilitates the construction of datasets for tasks like Text-to-SQL.

However, obtaining such corpora in practice remains highly challenging. The process typically requires labor-intensive manual annotation by domain experts, resulting in high costs in terms of time and resources. Meanwhile, publicly available SQL corpora are scarce, as real-world workloads are often restricted due to commercial sensitivity, or limited access to production environments. Therefore, automatic SQL generation has gained increasing attention [11, 64]. Prior research has explored a variety of techniques,

including (1) Template-based methods, which rely on hand-crafted query templates to produce interpretable yet structurally limited SQL instances [26, 37, 42]; (2) Grammar-guided random generation, which constructs queries by sampling from SQL grammar rules but often lacks semantic coherence [1, 46, 47]; and (3) Learning-based approaches, which employ machine learning models to generate queries in a data-driven manner [27, 32, 64]. However, existing methods often struggle to simultaneously achieve sufficient query diversity, schema-aware control, and cost-effective generation.

**Insufficient query diversity.** A lack of diversity means that the generated queries tend to follow similar structures, use the same operators, or involve limited combinations of joined tables. Template-based approaches such as TPC-DS [37] and Join Order Benchmark (JOB) [26] typically rely on fixed syntactic forms, which restrict their ability to produce varied query types. Learning-based methods such as TreeGAN [32] and LearnedSQLGen [64] exhibit strong tendencies to overfit to frequent query skeletons in training corpora, producing workloads clustered around recurrent table subsets and operator chains. As a result, these methods fail to capture the rich variety of queries that appear in real-world workloads. Improving query diversity would help cover more edge cases and make trained models more robust in practical database environments.

**Lack of schema-aware planning.** Effective SQL generation should account for schema heterogeneity by selectively targeting complex or underutilized tables. For instance, tables with many columns or dense foreign key connections are typically involved in more intricate query logic and should therefore receive more attention during workload construction. However, most existing methods treat all tables equally, applying uniform sampling that ignore these schema-level differences. Random generation methods such as SQL-Smith [46] sample syntax trees without considering schema semantics or historical usage patterns, leading to generating queries that either overfit frequent tables or neglect structurally rich regions of the schema. Although OmniSQL [27] incorporates schema context into its prompts, it does not explicitly guide table selection based on schema-level statistics. This lack of adaptive planning limits its ability to ensure balanced coverage or prioritize structurally informative parts of the schema.

**High generation cost.** Existing methods based on large language models (LLMs) often depend on either proprietary APIs or heavy-weight open-source models to achieve reliable quality, leading to high costs in API usage or deployment. For instance, OmniSQL [27] demonstrates the feasibility of large-scale SQL generation by using a mixture of locally deployed models, including those with over 30 billion parameters. This approach significantly increases deployment costs.

To address these limitations, we introduce SQL-Factory, a multi-agent framework designed for high-quality and large-scale SQL query generation. SQL-Factory leverages three collaborative agent teams: the Management Team, the Generation Team, and the Expansion Team, each responsible for a distinct aspect of the generation process. The **Generation Team** explores structurally diverse and expressive SQL patterns using a powerful large-scale language model, ensuring sufficient variety across query forms and structures. The **Expansion Team** complements this by scaling the workload through a lightweight and cost-effective local model, which generate semantically coherent variants based on representative

seed queries, significantly reducing reliance on expensive models. Finally, the **Management Team** evaluates query quality and schema-level distribution in real time, guiding scheduling decisions to ensure that structurally complex or underrepresented schema components are adequately covered. This division of responsibilities enables SQL-Factory to achieve a promising trade-off between query diversity, generation efficiency, and schema-level balance. Our contributions are summarized as follows:

- We propose **SQL-Factory**, a multi-agent SQL generation framework that introduces three specialized teams for generation, expansion, and management, enabling high-quality and large-scale SQL synthesis.
- We design a self-correction generation workflow that dynamically integrates structural exploration (Generation) and pattern exploitation (Expansion). The Management Team adaptively schedules different agents based on real-time evaluation of query quality.
- We generate over 300,000 broadly distributed SQL queries across four widely used benchmarks: TPC-DS, IMDB, BIRD, and SPIDER. Extensive experiments demonstrate that these queries exhibit significantly higher structural and semantic diversity compared to existing generation methods, and further improve performance on downstream tasks such as Text-to-SQL modeling and SQL clustering.

## 2 RELATED WORK

### 2.1 Data-Driven Intelligent Databases

The integration of data-driven artificial intelligence techniques, which leverage models trained on large-scale datasets, into database systems has emerged as a rapidly advancing research area. This paradigm explores the use of learning-based approaches, particularly large language models, to improve various database-related tasks such as configuration optimization, cardinality estimation, and Text-to-SQL translation.

One major area of research focuses on database configuration tuning. Various automatic tuning methods [30, 51, 63] have been proposed to recommend high-performance system configurations, thereby reducing manual effort and improving overall throughput. These methods typically rely on learning from historical workload logs to inform future configuration decisions.

Another important line of work focuses on converting natural language to SQL (NL2SQL). A wide range of methods [12, 14, 23, 43, 52, 59, 62] train neural models on aligned natural language and SQL query pairs, enabling users to express complex information needs in plain language and receive executable queries as output.

Query clustering is also an important task, which aims to group structurally or semantically similar SQL queries in order to enable log summarization, workload compression, and behavior profiling. Prior work has proposed a variety of approaches based on syntactic structures [25] and learned representations [49]. These methods typically require access to a diverse and well-structured corpus of queries to achieve meaningful clustering outcomes. However, such corpora are often unavailable or incomplete in real-world scenarios.

Beyond these applications, research in this domain have also been applied to cardinality estimation [31, 38, 53] and index selection [55, 61], where models benefit from the ability to generalize

across query patterns and data distributions. These advances reflect the broader trend of integrating deep learning methods into database system components.

Despite these advancements, most of them rely on the availability of high-quality SQL corpus for training or simulation. However, obtaining such a corpus is highly challenging in practice due to the labor-intensive and costly manual annotation by domain experts familiar with SQL and schema semantics, as well as the limited availability of real-world queries. This creates a growing need for automatic SQL generation frameworks.

## 2.2 SQL Generation

SQL generation aims to automatically synthesize large volumes of semantically meaningful and structurally diverse queries tailored to a specific database schema or system. This task plays an essential role in supporting downstream applications such as system benchmarking, robustness testing, and data-driven model training.

Early methods for SQL generation were primarily rule-based or template-driven. For example, industry-standard benchmarks such as TPC-DS [37], TPC-H [42] and Join Order Benchmark (JOB) [26] construct queries from pre-defined templates, offering a controllable and interpretable means of workload construction. Tenet[6] similarly adopts a template-driven approach, but introduces seed annotations and query types to synthesize more semantically diverse SQL queries. However, due to the limited number of templates, these methods often struggle to cover diverse query patterns and exhibit limited query randomness or structural variability. Random generation methods [1, 46, 47] attempt to increase diversity by randomly walking through SQL grammar rules or syntax trees. These methods increase lexical diversity but lack semantic grounding, often producing unexecutable or meaningless queries that fail to reflect realistic workload patterns.

Recent research has explored the use of machine learning and large language models for SQL generation. For example, TreeGAN [32] proposed a GAN-based framework that synthesizes SQL queries by imitating the structure of existing workloads. However, it requires large volumes of representative SQL data for training and struggles to generalize beyond seen patterns or schemas. LearnedSQLGen [64] introduces a reinforcement learning, aiming to synthesize queries that satisfy specific execution constraints learning approach that generates queries under execution constraints such as cardinality or cost. However, it often lacks structural diversity, as the model tends to exploit familiar query patterns that satisfy the constraints rather than exploring novel or complex structures. OmniSQL [27] adopts a large-scale data generation pipeline that synthesizes millions of queries using a set of locally employed models. While it achieves impressive scale, the framework relies heavily on high-cost models and lacks intermediate control during query synthesis. Without mechanisms to guide schema usage or enforce diversity, OmniSQL often produces redundant queries and cannot adapt generation focus in real time. These limitations reduce its efficiency, especially under resource constraints.

## 2.3 LLM-based Multi-Agent System

Large language models (LLMs) have demonstrated remarkable capabilities in natural language understanding and reasoning, enabling

them to perform a wide range of complex tasks with human-like proficiency. To organize these capabilities effectively, single-agent architectures typically equip an LLM with modular components such as Profile, Memory, Planning, and Action [57], which allow it to decompose tasks, retrieve relevant information, and simulate multi-step workflows. This design has proven effective in various domains, including dialogue generation, retrieval-augmented generation, and structured data interaction [19, 35, 54].

Building upon this foundation, recent research has extended LLM capabilities into collaborative multi-agent systems, where multiple specialized agents coordinate to complete tasks through dynamic communication and role-based reasoning. Multi-agent systems distribute cognitive responsibilities across distinct agents, each tailored for specific subgoals or expertise areas. Multi-agent LLM systems have shown promise in domains such as software engineering [8, 45], where agents representing roles like developer, tester, and reviewer to build and refine software artifacts. Similarly, in scientific experimentation [67] and medical diagnostics [50], these systems enable complex reasoning chains by leveraging role specialization and interactive deliberation among agents.

In the context of databases, LLM-based multi-agent approaches have recently emerged as a powerful solution for tasks requiring high modularity and semantic precision. For instance, in Text-to-SQL translation, multiple agents can collaboratively handle subtasks such as schema linking, query generation, and final SQL refinement, leading to more robust and accurate SQL outputs [52, 58]. In the realm of database anomaly diagnosis, systems like D-Bot [68] leverage several large language models to automatically analyze and identify root causes of anomalies by extracting knowledge from diagnostic documents. Additionally, ROMAS [20] introduces a role-based multi-agent framework for database monitoring and planning, featuring agents assigned to specific roles (planner, monitor, and worker) collaboratively facilitate dynamic task allocation, self-monitoring, and adaptive re-planning in complex data environments. These examples demonstrate the potential of multi-agent architectures to enhance reasoning modularity, improve decision accuracy, and integrate system-level feedback more effectively. Despite these advances, the application of LLM-based multi-agent systems to large-scale SQL generation remains underexplored.

## 3 SYSTEM OVERVIEW

This section introduces SQL-Factory, a schema-aware and cost-efficient multi-agent framework for automatically synthesizing large-scale and diverse SQL query workloads. As shown in Figure 2, SQL-Factory consists of three collaborative agent teams: (1) the **Management Team**, (2) the **Generation Team**, and (3) the **Expansion Team**. These teams work together to address three key challenges in SQL generation, including insufficient structural diversity, lack of schema-aware planning, and high generation cost. In addition, SQL-Factory maintains a centralized **SQL Pool** that stores all synthesized queries throughout the process.

**Generation Team.** The Generation Team is designed to maximize structural diversity of generated queries, tackling the problem of repetitive or overly templated query outputs. It comprises the Table Selection Agent and the Generation Agent. The Table Selection Agent realizes schema-aware planning by prioritizing structurally

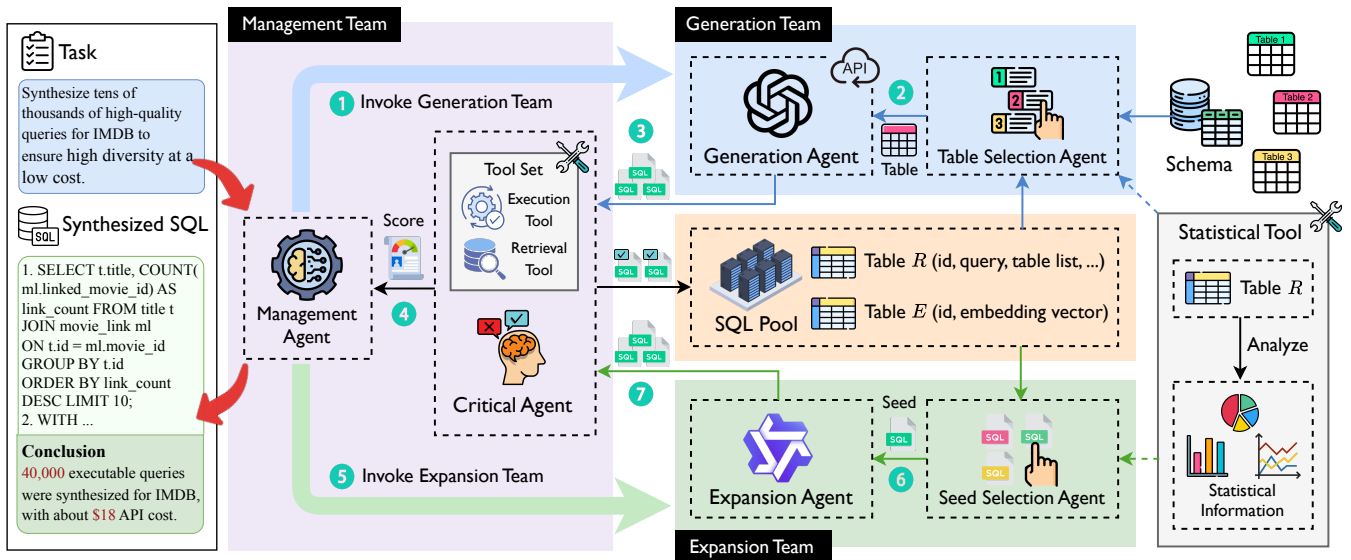


Figure 2: An overview of SQL-Factory.

complex tables or those with few synthesized queries. Then, the Generation Agent employs a powerful LLM to generate novel and semantically valid queries for selected tables. These queries often include complex structures such as nested subqueries, aggregations, and multi-table joins. By focusing on pattern exploration, the Generation Team ensures the introduction of high-diversity query structures that guide subsequent expansion.

**Expansion Team.** The Expansion Team tackles the cost-efficiency challenge by scaling up the number of queries using lightweight generation. It includes the Seed Selection Agent and the Expansion Agent. The Seed Selection Agent identifies high-quality queries that involve under-covered tables from the SQL Pool, and the Expansion Agent uses them as seed queries to synthesize new variants via controlled augmentation. The Expansion Team concentrates on scaling promising patterns to achieve high query volume with minimal computational overhead, enabling the system to achieve both depth and breadth in SQL generation.

**Management Team.** The Management Team serves as the central controller of SQL-Factory, responsible for coordinating these two teams and improving overall quality by filtering out invalid or overly similar queries. It plays a crucial role in maintaining system efficiency and adaptability, directly addressing the challenge of balancing structural diversity with generation cost. This team consists of two components: the Critical Agent and the Management Agent. The Critical Agent performs real-time quality assessments using an Execution Tool to check query validity and a Retrieval Tool to estimate semantic redundancy via hybrid similarity. Based on these evaluations, the Management Agent dynamically switches between exploration (Generation Team) and exploitation (Expansion Team) to balance novelty and efficiency.

**SQL Pool.** The SQL Pool is the centralized repository for all generated and expanded SQL queries. Implemented with a relational database such as MySQL, it is composed of two tables: Table  $R$ , which stores the detailed SQL query along with the list of referenced tables, and Table  $E$ , which maintains precomputed embedding vectors for similarity measurement and retrieval tasks.

## 4 METHODOLOGY

In this section, we provide a detailed methodology for SQL-Factory by describing the internal design and its key components, including the Management Team, Generation Team and Expansion Teams.

### 4.1 Management Team

The Management Team is responsible for orchestrating the entire SQL synthesis pipeline by coordinating interactions between the Generation Team and the Expansion Team. Its core objective is to maintain a balanced trade-off between structural diversity and data scalability while ensuring that only high-quality SQL queries are retained. This team consists of two components: the **Management Agent**, which governs the overall coordination of SQL synthesis, and the **Critical Agent**, which evaluates the quality of SQL queries from multiple perspectives.

**4.1.1 Management Agent.** The Management Agent governs the system state by dynamically selecting whether the Generation Team or Expansion Team should be invoked in each round. Its goal is to adaptively balance between generating structurally novel queries (exploration) and expanding promising query patterns based on existing queries (exploitation). This agent is implemented using a large language model, enabling flexible reasoning over both system statistics and quality metrics obtained from the Critical Agent.

Formally, the scheduling process is modeled as a discrete state transition function:

$$S_{t+1} = \text{Agent}_M(S_t, E_t), \quad (1)$$

where  $S_t \in [\text{GEN}, \text{EXP}]$  denotes the current system state (i.e., invoking the Generation or Expansion Team), and  $E_t$  represents the quality evaluation result provided by the Critical Agent at round  $t$ , including metrics such as query executability and average structural similarity. The policy function  $\text{Agent}_M(\cdot)$  is implemented via LLM-based inference, which selects the next state  $S_{t+1}$  based on statistical trends observed in the evaluation history.

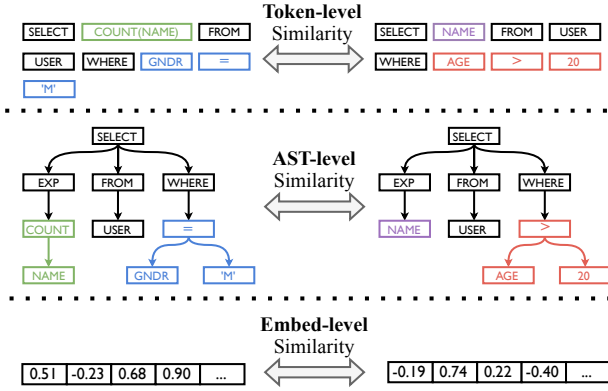


Figure 3: An overview of our hybrid similarity.

Specifically, the system begins in the exploration phase, where the Generation Team synthesizes queries for structurally complex tables or those with few synthesized queries. After a fixed number of generation cycles, control shifts to the Expansion Team to exploit learned patterns. During this exploitation phase, the Management Agent continuously monitors evaluation information from the Critical Agent. If query redundancy increases significantly or executability declines, it triggers a transition back to exploration. This alternating control mechanism allows SQL-Factory to incrementally grow a dataset that remains both diverse and scalable, while preventing the generation process from collapsing into repetitive or overly common query patterns. Compared to rule-based heuristics or reinforcement learning, our LLM-based scheduling leverages schema statistics and top-k query neighbors returned by the Critical Agent to assess diversity and redundancy. This enables the agent to make informed, context-sensitive decisions across benchmarks without the need for retraining or hand-crafted rules.

**4.1.2 Critical Agent.** The Critical Agent serves as the evaluation engine of SQL-Factory. Its primary function is to assess the quality of newly generated or expanded SQL queries on the fly before they are stored into the SQL Pool. This agent operates as a shared component used by both the Generation Team and the Expansion Team, and its evaluation result also plays a crucial role in guiding the Management Agent’s scheduling decisions.

Specifically, the Critical Agent comprises two modular components: the **Execution Tool**  $\text{Tool}_E$ , which verifies query executability, and the **Retrieval Tool**  $\text{Tool}_R$ , which using hybrid similarity to retrieve the most similar SQL queries. The evaluation process differs slightly depending on whether the query was produced by the Generation Team or the Expansion Team.

For SQL queries from the Generation Team, the powerful large model ensures a degree of novelty and exploratory coverage. Therefore, we focus primarily on syntactic and semantic validity. Given a set of generated queries  $Q_t = \{q_t^1, q_t^2, \dots, q_t^n\}$  at step  $t$ , the Critical Agent call the Execution Tool  $\text{Tool}_E$  to compile each query on the target database schema and reports a binary executability flag. The overall evaluation for this batch is defined as:

$$E_t = \text{Agent}_C(Q_t, \text{Tool}_E), \quad (2)$$

In contrast, queries from the Expansion Team are subject to stricter evaluation to mitigate semantic redundancy. Each query is

first passed through the Execution Tool as above, and then evaluated with the help of the Retrieval Tool. The Retrieval Tool encodes each query into a dense embedding vector using a pretrained language model. It then performs approximate nearest neighbor (ANN) retrieval to collect the top- $(k \times 20)$  candidates from the SQL Pool based on cosine similarity. From the retrieved candidate set, we select the top- $k$  most relevant queries for final evaluation using a **hybrid similarity metric**.

We introduce this hybrid metric to address the inherent limitations of embedding-based cosine similarity, which primarily captures surface-level semantic proximity. While effective for general clustering tasks, cosine similarity often fails to capture deeper structural or logical differences between queries that have similar tokens but differ significantly in their execution semantics or syntax trees. To produce a more robust and semantically grounded similarity measure, we integrate multiple perspectives into a unified scoring framework. An illustration of the hybrid similarity computation is provided in Figure 3. Formally, the similarity between a candidate query  $q$  and a historical query  $q' \in \text{SQL Pool}$  is computed as:

$$\text{Sim}(q, q') = \alpha \cdot \text{Sim}_{\text{Tok}} + \beta \cdot \text{Sim}_{\text{AST}} + \gamma \cdot \text{Sim}_{\text{Emb}}, \quad (3)$$

where  $\alpha$ ,  $\beta$ , and  $\gamma$  are the weights assigned to the three components. Each component here captures a different aspect of query similarity:

- **Token-level similarity**  $\text{Sim}_{\text{Tok}}$  is computed via normalized token to measure lexical alignment after token sorting.
- **AST-level similarity**  $\text{Sim}_{\text{AST}}$  evaluates structural similarity based on the abstract syntax trees (ASTs) of the two SQL queries, defined as:

$$\text{Sim}_{\text{AST}}(q, q') = 1 - \frac{\text{EditDist}(\text{AST}(q), \text{AST}(q'))}{\max(|\text{AST}(q)|, |\text{AST}(q')|)}, \quad (4)$$

where  $\text{EditDist}(\cdot)$  expresses the edit distance between AST trees, and  $|\cdot|$  represents the total number of nodes in the AST. This reflects the structural overlap between query patterns, including shared clauses and logical operators.

- **Embedding-level similarity**  $\text{Sim}_{\text{Emb}}$  is computed using cosine similarity between dense representations. It captures semantic alignment in latent space, abstracting away surface-level variations, defined as:

$$\text{Sim}_{\text{Emb}}(q, q') = \frac{\text{Emb}(q) \cdot \text{Emb}(q')}{\|\text{Emb}(q)\| \|\text{Emb}(q')\|}, \quad (5)$$

where  $\text{Emb}(q)$  denotes the vector representation of  $q$  encoded by the BERT, and  $\|\cdot\|$  represents the norm of the vector.

Given the distinct statistical behavior and discriminative power of these metrics, we assign the three components unequal weights, with  $\alpha = 0.6$ ,  $\beta = 0.3$ , and  $\gamma = 0.1$ , to reflect their differing variance ranges and discriminative effectiveness. In particular, token-level similarity exhibits the widest score dispersion and is thus most effective for detecting fine-grained redundancy. AST-level and embedding-level signals, while complementary, exhibit narrower and more concentrated distributions, making them more suitable as auxiliary signals.

By integrating the above three similarity metrics, this hybrid scoring function provides a more comprehensive assessment of query similarity. It is particularly effective in identifying subtle redundancies among expanded SQL queries.

Instruction	You are a database expert. Below, you are provided with some tables and their corresponding database schemas. Your task is to generate ten complex and diverse SQL queries for these tables.
Tips	The query you generate needs to meet the following conditions. 1. The generated query needs to cover as many tables as possible. 2. It is better if the generated query contains very complex operators and window functions, such as CASE WHEN, RANK, DENSE_RANK and others. 3. The generated queries should not be too similar to each other. 4. You'd better think carefully before giving an answer. You are very good at this and will definitely do well. .....
Schema Info	Table Schema: {schema_info} Foreign Key Dependency: {foreign_info}
Output Format	Each generated query must be output in the following format: <start-sql> [Query1 here] <end-sql> <start-sql> [Query2 here] <end-sql> .....

Figure 4: Generation Prompt Template.

Finally, the Critical Agent aggregates the execution validation results from  $\text{Tool}_E$  and retrieval results from  $\text{Tool}_R$  to produce the final evaluation score.

$$E_t = \text{Agent}_C(Q_t, \text{Tool}_E, \text{Tool}_R), \quad (6)$$

The final evaluation result  $E_t$  is then passed to the Management Agent to support real-time scheduling decisions.

## 4.2 Generation Team

The Generation Team is responsible for initiating new SQL patterns through schema-aware query construction. This team is activated when the system enters the exploration phase, typically after the Management Agent detects either insufficient diversity or excessive redundancy in the current SQL Pool. The primary goal of the Generation Team is to expand the search space by synthesizing novel and complex queries, thereby enriching the structural and logical variety of the dataset.

To achieve this, the Generation Team operates in a two-step pipeline. First, the **Table Selection Agent** analyzes schema-level statistics from the SQL Pool to identify underrepresented tables that have not been sufficiently queried. Then, the **Generation Agent** uses a large language model to construct diverse SQL queries over the selected tables. These queries are generated from scratch, without direct reference to existing SQL patterns, ensuring minimal bias and maximal novelty.

By strategically prioritizing structurally rich but underutilized schema components and leveraging the expressive power of LLMs, the Generation Team enables SQL-Factory to systematically explore the vast query space and discover new SQL structures that are unlikely to emerge from pattern-based expansion alone.

**4.2.1 Table Selection Agent.** The Table Selection Agent is responsible for identifying which tables should be prioritized for SQL generation during the exploration phase. Its goal is to maximize structural diversity by focusing on underrepresented or structurally rich tables in the database. This agent forms the first step in the

generation pipeline and directly influences the quality and novelty of queries synthesized by the subsequent Generation Agent.

The agent operates over two sources of information: (1) the full schema of the target database, including all table names and their respective column types; and (2) a statistical summary of SQL usage derived from the SQL Pool, such as the number of queries for each table. These metrics are extracted and structured via a dedicated **Statistical Tool**  $\text{Tool}_S$ , which computes table-level query coverage.

Formally, let  $\mathcal{T} = \{T_1, T_2, \dots, T_m\}$  denote the complete set of tables in the current database. The full input to the Table Selection Agent is then:

$$\mathcal{S} = \{ \langle \text{Schema}(T_i), \text{Tool}_S(\text{SQL Pool}, T_i) \rangle \mid T_i \in \mathcal{T} \}, \quad (7)$$

The LLM is then prompted with  $\mathcal{S}$  to select a subset of tables  $\mathcal{T}' \subset \mathcal{T}$  that are sparsely covered. Formally, the output of the Table Selection Agent is defined as:

$$\mathcal{T}' = \text{Agent}_T(\mathcal{S}), \quad \text{where } \mathcal{T}' \subset \mathcal{T}. \quad (8)$$

This agent operates over the full schema, and the underlying LLM can implicitly consider factors such as the number of columns, referential degree, and depth of the referential tree when selecting tables. These schema-level indicators highlight structural importance and connectivity. At the same time, the agent also accounts for query statistics, giving higher priority to tables that are currently underrepresented in the SQL Pool. In this way, both structurally complex tables and those with relatively few generated queries are emphasized. By integrating schema-level complexity with query-level statistics, the Table Selection Agent allows informed decision-making that balances exploration of new structural patterns with targeted improvement of under-sampled schema regions. This design enables dynamic adaptation as the SQL Pool evolves, ensuring diversity as well as representativeness in the generated queries.

**4.2.2 Generation Agent.** The Generation Agent is responsible for synthesizing new SQL queries in a schema-aware and open-ended manner. It takes as input the set of tables  $\mathcal{T}'$  selected by the Table Selection Agent and aims to produce queries that are structurally diverse. This process constitutes the core exploration step in SQL-Factory, enabling the discovery of novel query patterns that cannot be easily obtained by simply modifying existing SQL queries.

The agent is implemented using a powerful, instruction-tuned large language model, which is prompted with both schema information for  $\mathcal{T}'$  and high-level generation instructions, as shown in Figure 4. Rather than treating each table independently, the model considers the joint schema of all selected tables and is encouraged to synthesize complex queries that include multi-table joins, nested subqueries, aggregations, and window functions. Besides, to guide the model effectively, the schema information included in the prompt is carefully constructed to ensure both syntactic validity and semantic meaningfulness. In particular, we also incorporate explicit foreign key relationships, allowing the model to generate meaningful join conditions that align with actual schema constraints. For each selected table, we encode not only its column names and data types, but also relevant content-aware hints to reduce the likelihood of generating vacuous or logically invalid queries. Specifically, for columns of enumerated types (e.g., status, category), we provide up to five representative values from the domain to help the model infer plausible filter conditions or join

Instruction	You are a database expert, skilled in writing various SQL statements. Below, you are provided with three SQL queries. Your task is to analyze the schema-related information and then generate ten diverse SQL queries based on them.
Tips	There are some important tips you need to remember: 1. Using schema information that does not exist in the given query is not allowed, don't try to create or modify the name of table or column, even if the origin column is unreasonable. 2. The generated queries should be diverse enough and not similar to each other. 3. You'd better think step by step carefully before giving an answer. You are very good at this and will definitely do well. .....
Reference Query	Queries: {query list}
Output Format	Each generated query must be output in the following format: <start-sql> [Query1 here] <end-sql> <start-sql> [Query2 here] <end-sql> .....

Figure 5: Expansion Prompt Template.

constraints. For numeric or date type columns (e.g., int, float, date), we include its minimum and maximum values in the database, offering statistical context for range-based predicates (e.g., price > 100). For free-text or string columns, we randomly sample up to five values as textual examples, assisting the model in generate meaningful LIKE conditions or equality conditions. These content-aware augmentations are inserted inline with the schema description in a structured format, forming an enriched prompt that provides the model with both structural knowledge and semantic information. This design improves contextual relevance and guides the model away from implausible or extreme filter conditions, enhancing overall semantic plausibility. This is particularly beneficial for tables with sparse usage history or complex attributes. Finally, the output of the Generation Agent are summarized as:

$$Q_{\text{gen}} = \text{Agent}_G(\mathcal{T}') = \{q_1, q_2, \dots, q_k\}. \quad (9)$$

The generation is carried out in an open-ended fashion, without referencing queries from the SQL Pool. To ensure structural richness, prompts are explicitly designed to avoid trivial query types such as single-table scans or simple projections. This setup allows the Generation Agent to push the boundaries of the query space and introduce SQL patterns that improve structural diversity. All generated queries  $Q_{\text{gen}}$  are forwarded to the Critical Agent for executability validation. Only those that compile successfully are inserted into the SQL Pool. In this way, the Generation Agent supports horizontal dataset expansion, discovering new structural dimensions of the underlying database schema.

### 4.3 Expansion Team

The Expansion Team is designed to scale up the SQL dataset by exploiting existing query structures through controlled augmentation. This team is activated during the exploitation phase, particularly after the Management Agent determines that the diversity of the SQL Pool has been sufficiently enriched and the system should focus on increasing query volume. Unlike the Generation Team, which relies on a large-scale language model to generate structurally novel queries from scratch, the Expansion Team utilizes a

lightweight model to synthesize new SQL queries by learning from high-quality seeds already present in the SQL Pool.

The expansion process also follows a two-step workflow. First, the **Seed Selection Agent** identifies representative queries from the SQL Pool that involve under-covered tables or patterns. Then, the **Expansion Agent** learns from these queries and generates structurally similar but lexically and semantically varied SQL queries that preserve executability. Since this phase focuses on local exploration around known patterns, it enables SQL-Factory to rapidly increase the number of high-quality queries while maintaining system efficiency. Together, these two agents complements the structural exploration performed by the Generation Team and allows SQL-Factory to produce a large volume of diverse and executable queries at scale with low cost.

**4.3.1 Seed Selection Agent.** The Seed Selection Agent plays a similar role to the Table Selection Agent in the Generation Team, as both are responsible for identifying underrepresented schema components for further query synthesis. However, rather than initiating query generation from scratch, the Seed Selection Agent focuses on retrieving representative queries from the existing SQL Pool to serve as seeds for the Expansion Agent. These seed queries act as structural blueprints, enabling the system to scale existing patterns while preserving semantic integrity.

To select appropriate seeds, the agent first determines a set of underutilized tables  $\mathcal{T}' \subseteq \mathcal{T}$  by analyzing query distribution statistics via a shared statistical module  $\text{Tools}_S$ . It then search the SQL Pool to collect a small set of seed queries  $Q_{\text{seed}}$ , each of which must reference at least one table from  $\mathcal{T}'$ . These queries are selected not only based on table inclusion but also for their structural validity and potential for lexical or syntactic variation. This approach allows the Expansion Agent to operate in a pattern-aware and data-driven manner, leveraging past structures to efficiently generate new and diverse SQL variants. Formally, the the final output of Seed Selection Agent is:

$$Q_{\text{seed}} = \text{Agents}_S(\mathcal{T}') = \{q_i \in Q_{\text{pool}} \mid \text{Tables}(q_i) \cap \mathcal{T}' \neq \emptyset\}. \quad (10)$$

By identifying high-quality yet underutilized queries, the Seed Selection Agent enables the Expansion Team to efficiently enrich the dataset without redundant exploration.

**4.3.2 Expansion Agent.** The Expansion Agent generates new SQL queries by leveraging representative seed queries retrieved by the Seed Selection Agent. Unlike the Generation Agent, which synthesizes queries from scratch based on schema information, the Expansion Agent focuses on local exploration. Its goal is to generate syntactically diverse and executable queries by learning from the structure or intent of seed queries. This agent enables the system to scale up query volume without incurring the computational cost of large-scale model inference.

Specifically, this agent operates on the set of seed queries, which were selected based on their relevance to underrepresented tables. It uses a lightweight model to derive multiple variants from these seed queries, employing techniques such as structural rewriting (e.g., rewriting subqueries as CTEs), substructure recombination (e.g., merging partial queries sharing common tables), lexical variation (e.g., alias renaming or clause reordering), and logical variation (e.g., modifying filter thresholds or switching conjunctions). These

transformations are guided implicitly by the prompt rather than fixed rules, enabling flexible and semantically rich local exploration.

As shown in Figure 5, each prompt includes a small set of seed SQL queries and task-specific instructions. These instructions encourage the model to generate structurally diverse and executable variants, while inferring schema constraints implicitly from the provided queries. Formally, given a set of seed queries  $Q_{\text{seed}}$ , the Expansion Agent generates a corresponding set of SQL variants  $Q_{\text{exp}}$  as follows:

$$Q_{\text{exp}} = \text{Agent}_E(Q_{\text{seed}}). \quad (11)$$

After expansion, the resulting queries  $Q_{\text{exp}}$  are passed to the Critical Agent. In this phase, queries are checked for executability and evaluated for semantic redundancy using retrieval-based similarity assessment. Only queries satisfying both executability and novelty criteria are retained in the SQL Pool. Through this process, the Expansion Agent achieves vertical dataset growth by increasing the density of high-quality query patterns within the local schema space, enhancing the dataset’s expressiveness and utility for downstream tasks. While the Generation Agent explores new structural forms, the Expansion Agent intensifies local pattern regions, ensuring both diversity and volume.

## 5 EXPERIMENTS

In this section, we conduct extensive experiments to demonstrate the efficiency of the SQL-Factory framework from multiple perspectives, including detailed analysis on query quality, and its effectiveness on downstream tasks.

### 5.1 Experiment Setup

**5.1.1 Environment.** We conduct the experiments on one GPU servers, equipped with 4 NVIDIA RTX-4090-24GB GPUs, an Intel(R) Xeon(R) Silver 4316 CPU, and 256GB of RAM. For LLM training, we utilize PyTorch 2.3.0, DeepSpeed 0.14.0 and LLaMA-Factory 0.8.4, while vLLM 0.4.3 is employed for LLM inference.

**5.1.2 Benchmarks.** To evaluate the effectiveness and generality of SQL-Factory, we select four typical database benchmarks for synthesizing queries. (1) **TPC-DS** [37] is an industry-standard benchmark for decision support systems, featuring a complex schema with 24 tables and various star-join patterns. (2) **IMDB** [33] is a movie-related database, which includes 21 tables on movies and actors. Its real-world structure and moderate size make it suitable for generating realistic join-heavy queries. (3) **SPIDER** [60] is a large-scale text-to-SQL dataset consisting of several databases with highly diverse and normalized schemas. (4) **BIRD** [29] is a recent benchmark introduced to evaluate database-instruction reasoning in realistic and complex environments. Across all benchmarks, we parse the schema into structured representations including table names, column types, and foreign key relations. These are used as inputs for SQL-Factory’s Generation and Expansion workflows.

Besides, To evaluate the effectiveness of the generated SQL queries in improving representation learning, we adopt three widely-used query clustering benchmarks: **IIT Bombay** [7], **UB Exam**, and **PocketData** [22]. These datasets contain real-world SQL queries collected from academic and industry environments and their manually annotated cluster labels.

**5.1.3 Baseline.** We organize our baselines into three categories to cover different stages of evaluation: SQL generation, Text-to-SQL, and SQL query clustering.

**SQL Generation Baselines.** To compare the quality and diversity of generated SQL queries, we consider both random generation methods and learning-based generation methods. Specifically, we include SQLSmith [46], a grammar-driven query generator that produces random SQLs, and two representative learning-based methods: LearnedSQLGen [64], which learns to generate queries via reinforcement learning, and OmniSQL [27], which leverages large language models to synthesize large-scale queries. Additionally, we exclude template-based methods [26, 37, 42] from our comparison, as they rely heavily on predefined templates and are not suitable for generating massive and structurally diverse SQL workloads.

**Text-to-SQL Baselines.** To assess the effectiveness of the generated training data, we evaluate a series of commonly used models, including Qwen2.5-Coder-1.5B/3B [21], LLaMA3.2-1B/3B [2], and Deepseek-Coder-1.3B [18], under three training settings: (1) the original model, (2) fine-tuning using the official benchmark training data (e.g., Spider or BIRD), and (3) fine-tuning using the data produced by our SQL generation pipeline.

**Query Clustering Baselines.** For the SQL clustering task, we evaluate several established similarity models. These include statistical and heuristic approaches such as Aouiche [4], Aligon [3], and Makiyama [34], the pretrained model-based PreQR, and two general-purpose language models: BERT-Base and BERT-Large [10].

**5.1.4 Implementation Details.** We utilize different large language models and configurations to implement these agents in SQL-Factory. Each module is configured to balance quality and efficiency.

**Generation Agent** We adopt GPT-4o [39] as the backbone model for the Generation Agent, due to its strong capabilities in zero-shot reasoning and diverse SQL synthesis. The model is accessed through the OpenAI API, with generation configured using a temperature of 0.8 and top- $p$  sampling for better output diversity.

**Expansion Agent.** For the Expansion Agent, we utilize Qwen2.5-Coder-14B [21], a locally deployed open-source model optimized for code and SQL-related instruction following capability. It is served using vLLM to enable rapid query expansion. It is also use a temperature of 0.8 for diversity.

**Other Agents.** The Management Agent, Critical Agent, Table Selection Agent, and Seed Selection Agent also reused Qwen2.5-Coder-14B to reduce deployment overhead. Here we use a lower temperature of 0.3, aiming for more deterministic and stable reasoning behavior without sacrificing overall coherence.

### 5.2 Quality Analysis of Generated SQL

We apply SQL-Factory to synthesize large-scale SQL queries across four benchmarks. Each benchmark varies in schema size and complexity, enabling us to test the scalability and adaptability of our multi-agent generation framework.

**5.2.1 Overview of Generated Queries.** We analyze the generated SQL queries across several dimensions, including join depth, predicate usage, subquery nesting, aggregate usage, and token length. These metrics serve as proxies for expressiveness, structural diversity, and alignment with real-world query patterns.




Figure 6: Overview of SQL queries synthesized by SQL-Factory.

Table 1: The statistics of the generated SQL.

Benchmark	# DB	# Table	# Col/DB	# Queries	Price (\$)	Time (h)
TPC-DS	1	24	449	50,000	51.78	56.5
IMDB	1	21	108	40,000	18.47	48.2
SPIDER	69	795	26.82	130,000	42.16	233.3
BIRD	146	524	54.56	160,000	67.50	283.1

Figure 6(a) illustrates the distribution of join operations per query. Notably, 78.6% of all queries involve at least one join, and 30% contain two or more join. This indicates that SQL-Factory effectively generates multi-table queries, showcasing its capacity for relational reasoning and schema-aware generation. Figure 6(b) analyzes the number of predicates in queries. More than 58% of queries contain two or more predicates, with 9.8% having four or more. This confirms that SQL-Factory not only generates structurally rich queries but also embeds fine-grained filtering logic, which is critical for supporting meaningful data retrieval and realistic workload simulation. Besides, Figure 6(c) and Figure 6(d) reveal substantial use of nested subqueries and aggregation, confirming the system’s capability to produce advanced constructs. Furthermore, Figure 6(e) reports the distribution of token lengths. The x-axis represents the number of tokens for each query, while the y-axis shows the corresponding frequency. The distribution reveals a broad range of query lengths, confirming that SQL-Factory does not produce templated or overly uniform queries, but rather supports diverse query formulations at varying levels of complexity. In addition, we analyze the proportion of queries that return empty result sets. Only 15.34% of the queries fall into this category, suggesting that most generated queries are not only syntactically valid but also semantically meaningful.

Overall, these results confirm that SQL-Factory produces syntactically diverse and structurally expressive SQL queries, which is essential for training robust models in downstream applications.

**5.2.2 Cost Analysis of SQL Generation.** Table 1 summarizes the key statistics of SQL queries generated by SQL-Factory across four benchmarks. These include the number of databases, tables, average number of columns per database, the total number of generated queries, and the corresponding cost of API. Across these benchmarks, we observe that the total number of generated queries generally correlates with the schema’s complexity and scale. For example, the largest numbers of queries are generated for BIRD (160,000) and Spider (130,000), which contain the most tables and databases. In contrast, IMDB yields 40,000 SQLs due to its simpler, single-schema structure. This adaptive generation behavior emerges naturally from our multi-agent framework, which allocates more generation

Table 2: Comparison with Existing Query Generation Methods in Terms of Diversity.

Method	Hybrid Similarity	Vendi Score SQLEncoder	Vendi Score SimCSE
SQLSmith	<b>0.571</b>	2.94	2.35
LearnedGen	0.863	2.33	2.14
OmniSQL	0.716	1.97	5.58
SQL-Factory	0.637	<b>4.10</b>	<b>7.94</b>

rounds and expansions in response to richer or more varied schema coverage requirements.

Despite the varying generation needs, SQL-Factory maintains high efficiency. The total cost across all four benchmarks remains under \$200. Notably, if the entire workload were generated solely by GPT-4o without the Expansion Team, the API cost would approach \$1000, highlighting the substantial cost savings enabled by our hybrid design. Besides, generating the first 10,000 queries typically takes around 7 to 11 hours, depending on schema complexity and prompt length. As the number of existing queries grows, operations such as statistical analysis and hybrid similarity computation introduce additional overhead, leading to a gradual slowdown. Nevertheless, the overall throughput remains significantly higher than manual query construction. Together, these results demonstrate that SQL-Factory is capable of producing SQL workloads at scale, while maintaining both schema adaptivity and cost-effectiveness.

**5.2.3 Comparative Analysis on Query Diversity.** We conduct a comparative study to assess the query diversity achieved by SQL-Factory against three representative SQL generation methods: (1) SQL-Smith [46], a grammar-driven random query generator that walks through SQL syntax trees; (2) LearnedSQLGen [64], which employs a reinforcement learning framework guided by cost-based feedback to generate queries through a finite-state abstraction; and (3) OmniSQL [27], a LLM-based framework that leverages prompt engineering to generate millions of SQL queries.

For a fair comparison, we generate 10,000 SQL queries for the TPC-DS benchmark using each method. For OmniSQL, we adopt the original prompts provided by the authors and use both GPT-4o and Qwen2.5-Coder-14B for query generation. The proportion of queries generated by each model is kept consistent with the generation and expansion ratio used in SQL-Factory. We evaluate the diversity of generated queries across multiple evaluation metrics, including our hybrid similarity and vendi score.

**Table 3: Comparison with single-agent baselines.**



Method	Time	Price	Hybrid Similarity	Vendi Score SQLEncoder	Vendi Score SimCSE
GPT-4o	9.2h	104.6\$	0.672	3.84	4.44
Qwen-14B	23.7h	-	0.733	2.70	3.88
SQL-Factory	11.4h	15.22\$	<b>0.637</b>	<b>4.10</b>	<b>7.94</b>

**Hybrid Similarity.** We use our proposed hybrid similarity metric as described in Section 4.1.2 to measure the average pairwise similarity between queries. As shown in Table 2, SQL-Factory achieves a balanced similarity score of 0.637, significantly lower than LearnedSQLGen (0.863) and OmniSQL (0.688), suggesting that it avoids over-repetition while maintaining meaningful structure. Besides, the low similarity score for SQLSmith is due to its excessive syntactic randomness. It was originally designed for database stress testing, focusing on exposing engine-level bugs rather than generating semantically coherent queries. While its queries increase lexical variance, they often contain irregular constructs and unrealistic patterns that do not align with practical workloads.

**Vendi Score.** The Vendi Score [13] is a recently proposed metric for evaluating diversity based on the exponential of the entropy of the eigenvalues of a similarity matrix computed from the data. It does not rely on predefined labels or reference distributions, making it especially suitable for open-ended generation tasks. We utilize Vendi Score to assess the diversity of our generated datasets with two embedding backbones: SimCSE [15], a general-purpose sentence encoder trained via contrastive learning, and SQL-Encoder [44], a specialized model pretrained for SQL embeddings. Higher Vendi scores indicate greater overall diversity. As shown in Table 2, SQL-Factory achieves the highest Vendi Score under both encoders. In contrast, SQLSmith demonstrates low similarity but also fails to produce coherent queries. LearnedSQLGen suffers from high redundancy due to its pattern imitation and optimization-driven training. Besides, OmniSQL is limited by its monolithic generation strategy.

**5.2.4 Ablation Study on Agent Decomposition.** To evaluate the benefits of agent decomposition, we compare SQL-Factory with two single-agent baselines on the TPC-DS benchmark: (1) direct generation using GPT-4o, and (2) direct generation using Qwen2.5-Coder-14B. Each method generates 10,000 queries. As shown in Table 3, SQL-Factory achieves the lowest hybrid similarity and highest Vendi Scores, indicating better structural and semantic diversity. Unlike single-agent approaches, SQL-Factory combines seed generation with targeted expansion, avoiding mode collapse and improving output quality. In addition, SQL-Factory is more cost-efficient, with a significantly lower API cost and moderate time consumption, demonstrating better scalability for large-scale query generation.

**5.2.5 Schema-Aware Query Allocation Analysis.** To evaluate SQL-Factory’s ability to support schema-aware query planning, we compare it with OmniSQL [27] using two metrics: the number of queries per table and each table’s structural complexity. We define table complexity following established schema complexity metrics in database engineering literature [24, 40, 41]. Specifically, we incorporate three widely recognized structural factors: the number of attributes (NA), the referential degree (RD), and the depth of the

**Figure 7: Query and Complexity Distributions in OmniSQL.****Figure 8: Query and Complexity Distributions in SQL-Factory.**

referential tree (DRT). The final complexity is calculated as:

$$\text{Complexity}(T) = \text{NA}(T) + 2 \text{RD}(T) + \text{DRT}(T). \quad (12)$$

This formulation captures both local structural richness and broader relational connectivity. Figure 7 and Figure 8 illustrate the resulting complexity distributions for OmniSQL and SQL-Factory, respectively. Figure 7(a) shows that exhibits a highly skewed distribution, where a small number of tables dominate the workload while most tables receive very few queries. However, as shown in Figure 7(b), the distribution of table complexity in the same benchmark is much more balanced, with the majority of tables falling in a moderate complexity range. This misalignment suggests that OmniSQL does not account for schema structure when allocating queries, which may result in under-representation of complex tables. In contrast, SQL-Factory produces a query distribution that more closely matches schema complexity. Figure 8(a) shows that most tables receive a moderate number of queries, and this pattern aligns well with the table complexity distribution in Figure 8(b), where moderately complex tables are also the most common. This consistency indicates that SQL-Factory can allocate queries in a way that reflects the underlying schema properties, supporting schema-aware planning and improving the quality of the workload.

**5.2.6 Generation Stopping Criteria.** Although SQL-Factory can generate an unlimited number of SQL queries for any benchmark, it is also important to ensure high overall data quality. To this end, we employ a termination policy that avoids excessive accumulation of redundant or low-diversity queries. Specifically, we monitor two key metrics: (1) the **expansion ratio**, defined as the proportion of queries in the SQL Pool that originate from the Expansion Team, and (2) the **average hybrid similarity**, which reflects overall similarity.

As illustrated in Figure 9, the system is initialized by the Generation Team, which produces a small number of diverse seed queries. After this warm-up phase, the Expansion Team begins to take over, resulting in a rise in both the expansion ratio and similarity across the SQL Pool. This indicates that a growing number of queries are

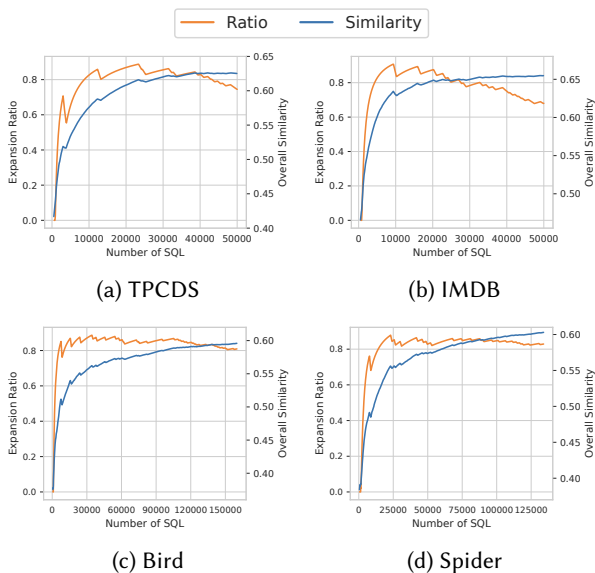


Figure 9: Evolution of Expansion Ratio and Overall Similarity as the Number of Synthesized SQL Queries Increases.

being synthesized based on previously explored patterns. To maintain overall query diversity, the Management Agent periodically switches back to the Generation Team. These transitions manifest as synchronized dips in both expansion ratio and similarity on the curve, reflecting the injection of more novel queries. This alternating pattern of generation and expansion continues iteratively, with diminishing gains over time. Once the Expansion Team becomes dominant and the similarity curve shows a sustained upward trend, it indicates that the SQL Pool has reached semantic saturation, prompting us to terminate the synthesis process.

For example, after synthesizing approximately 50,000 queries for TPC-DS, the Expansion Team begins to dominate the process. Specifically, the expansion ratio rapidly decrease while the average similarity among queries steadily rising. Continuing the process beyond this point leads to diminishing returns in diversity and semantic coverage. Therefore, we empirically adopt this saturation point as the stopping criterion for each benchmark.

### 5.3 Text-to-SQL Task

To further validate the practical value of the SQL-Factory framework, we evaluate its effectiveness in improving performance on downstream Text-to-SQL tasks.

**5.3.1 Dataset Construction.** Since SQL-Factory is designed to synthesize executable queries for a given database schema from scratch, applying it to Text-to-SQL tasks requires transforming these SQLs into natural language (NL) descriptions. This enables the construction of  $\langle \text{NL}, \text{SQL} \rangle$  training pairs for supervised learning.

Instead of generating SQL from NL, which often suffers from semantic mismatch and low exact match accuracy, we adopt a reverse formulation: converting SQL to NL. Recent studies [17, 23, 56, 65] have shown that LLMs are more reliable in generating faithful natural language descriptions for given SQL, making SQL-to-NL a robust and efficient method for constructing high-quality

Table 4: Results on Text-to-SQL Task (%). "+Official-Data" refers to fine-tuning using the original training sets provided by the benchmarks. "+SQL-Factory" refers to fine-tuning using training data generated by our SQL-Factory framework.

Method	Spider		Bird	
	Executable	Accuracy	Executable	Accuracy
<b>Qwen-1.5B</b>	85.7	69.2	57.6	28.6
+Official-Data	91.9 ( $\uparrow 6.2$ )	73.2 ( $\uparrow 4.0$ )	79.1 ( $\uparrow 21.5$ )	39.2 ( $\uparrow 10.6$ )
+SQL-Factory	<b>97.2 (<math>\uparrow 11.5</math>)</b>	<b>76.5 (<math>\uparrow 7.3</math>)</b>	<b>82.1 (<math>\uparrow 24.5</math>)</b>	<b>40.9 (<math>\uparrow 12.3</math>)</b>
<b>Qwen-3B</b>	92.8	76.2	68.8	37.1
+Official-Data	94.8 ( $\uparrow 2.0$ )	77.8 ( $\uparrow 1.6$ )	<b>88.7 (<math>\uparrow 19.9</math>)</b>	45.7 ( $\uparrow 8.6$ )
+SQL-Factory	<b>97.3 (<math>\uparrow 4.5</math>)</b>	<b>79.7 (<math>\uparrow 3.5</math>)</b>	88.1 ( $\uparrow 19.3$ )	<b>46.1 (<math>\uparrow 9.0</math>)</b>
<b>LLama-1B</b>	73.6	41.3	43.3	11.3
+Official-Data	80.1 ( $\uparrow 6.5$ )	60.3 ( $\uparrow 19.0$ )	62.5 ( $\uparrow 19.2$ )	26.0 ( $\uparrow 14.7$ )
+SQL-Factory	<b>94.3 (<math>\uparrow 20.7</math>)</b>	<b>70.6 (<math>\uparrow 29.3</math>)</b>	<b>71.4 (<math>\uparrow 28.1</math>)</b>	<b>31.4 (<math>\uparrow 20.1</math>)</b>
<b>LLama-3B</b>	89.0	60.4	73.9	30.1
+Official-Data	91.8 ( $\uparrow 2.8$ )	72.2 (11.8)	82.0 ( $\uparrow 8.1$ )	38.9 ( $\uparrow 8.8$ )
+SQL-Factory	<b>98.7 (<math>\uparrow 9.7</math>)</b>	<b>76.8 (16.4)</b>	<b>84.7 (<math>\uparrow 10.8</math>)</b>	<b>42.1 (<math>\uparrow 12.0</math>)</b>
<b>Deepseek-1.3B</b>	87.2	50.2	63.3	20.3
+Official-Data	88.7 ( $\uparrow 1.5$ )	70.7 ( $\uparrow 20.5$ )	<b>80.3 (<math>\uparrow 17.0</math>)</b>	36.9 ( $\uparrow 16.6$ )
+SQL-Factory	<b>98.3 (<math>\uparrow 11.1</math>)</b>	<b>75.7 (<math>\uparrow 25.5</math>)</b>	79.8 ( $\uparrow 16.5$ )	<b>38.3 (<math>\uparrow 18.0</math>)</b>

datasets. To this end, we develop a multi-stage pipeline for SQL-to-NL conversion with consistency checking and optional SQL refinement. The steps are as follows:

**(1) Similar Example Retrieval.** To improve translation quality and ensure semantic alignment, we first retrieve top- $k$  similar examples (with  $k=5$  in the experiment) from training set of BIRD or Spider using Retrieval Tool proposed in Section 4.1.2. These examples are used to construct the in-context prompt to teach model how to write fluent, semantically faithful questions from SQL. This retrieval-augmented generation ensures structural and linguistic alignment between training and inference distributions.

**(2) SQL-to-NL Conversion.** We then adopt a prompt-based method to translate SQL queries into natural language using a powerful instruction-tuned language model (Qwen2.5-Coder-14B[21]). Given the database schema, SQL query and a few-shot template, the model generates a corresponding natural language question that accurately reflects the intent of the SQL. Additionally, to ensure consistency with the BIRD dataset format, each generated NL question is accompanied by an *external knowledge explanation*.

**(3) Semantic Consistency Verification.** To ensure the correctness of the constructed  $\langle \text{NL}, \text{SQL} \rangle$  pairs, we introduce an LLM-based verifier to perform semantic consistency checking. This component filters out instances where the NL question does not accurately reflect the intent or logic of the query. While LLMs are not perfect in complex generative tasks like Text-to-SQL, recent work [9, 43, 52] shows they are highly reliable for binary classification and entailment-style tasks, supporting our choice of an LLM-based verifier. Specifically, the verifier takes as input the  $\langle \text{NL}, \text{SQL} \rangle$  pair, and is prompted to return a binary decision indicating whether the pair is semantically aligned. To further improve reasoning quality, we adopt a Chain-of-Thought (CoT) prompting strategy, where the model generates reasoning steps in a step-by-step manner before producing the final decision. In addition, we conducted a human-in-the-loop validation by randomly sampling 1,000 pairs and recruited five PhD-level STEM students with SQL experience to perform

**Table 5: Performance on Query Clustering Task.**

Method	IIT Bombay	UB Exam	PocketData
<b>Aouiche</b>	0.577	0.923	0.893
<b>Aligon</b>	0.535	0.799	0.898
<b>Makiyama</b>	0.665	0.897	0.879
<b>PreQR</b>	0.387	0.622	0.710
<b>BERT-Base</b>	0.524	0.769	0.773
+SQL-Factory	0.376 ( $\downarrow$ 0.148)	0.558 ( $\downarrow$ 0.211)	0.656 ( $\downarrow$ 0.117)
<b>BERT-Large</b>	0.637	0.846	0.777
+SQL-Factory	<b>0.350 (<math>\downarrow</math>0.287)</b>	<b>0.549 (<math>\downarrow</math>0.297)</b>	<b>0.638 (<math>\downarrow</math>0.139)</b>

cross-annotation. Results show that only 4.6% of the samples were judged to be semantically inconsistent, indicating that the final dataset maintains high semantic fidelity and practical reliability.

This multi-stage pipeline enables us to automatically transform SQL-only corpora generated by SQL-Factory into a large-scale, high-quality Text-to-SQL training set, effectively bridging the gap between unsupervised SQL generation and downstream tasks.

**5.3.2 Model Training.** We fine-tune all the backbone models mentioned above, including Qwen2.5-Coder-1.5B/3B, LLaMA3.2-1B/3B, and Deepseek-Coder-1.3B. Training is performed using the AdamW optimizer with hyperparameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.95$ , and  $\epsilon = 1 \times 10^{-8}$ . We use a peak learning rate of  $3 \times 10^{-6}$ , linearly warmed up over the first 5% of steps, followed by cosine learning rate decay. Each model is trained for 3 epochs with a batch size of 256. All other settings follow the default configuration of LLaMA-Factory [66].

**5.3.3 Result and Analysis.** To ensure fairness, we adopt the official prompt templates for each model. Each model is evaluated on the Spider and BIRD benchmarks, using both executability and execution accuracy as metrics. The results in Table 4 show that across all model scales and architectures, fine-tuning with SQL-Factory consistently improves performance on both benchmarks.

On Spider, Qwen2.5-Coder-1.5B improves in accuracy from 69.2% to 76.5%, while on BIRD, Deepseek-Coder-1.3B and LLaMA3.2-1B achieve accuracy gains exceeding 15–20 points. Similar trends are observed for executability. These gains are larger than those obtained from official datasets, indicating that SQL-Factory produces training data that rivals or surpasses expert-curated sets. These results demonstrate that data generated by SQL-Factory significantly boosts model performance. The improvements are particularly evident in structurally complex cases, indicating the effectiveness of our schema-aware generation and expansion process.

## 5.4 Query Clustering Task

To further evaluate the utility of the SQL queries generated by our SQL-Factory, we investigate their effectiveness in improving SQL representation for query clustering tasks.

We fine-tune both BERT-Base and BERT-Large models using all SQL queries synthesized by SQL-Factory, aiming to enhance their capability in capturing SQL-level semantics. The training process employs contrastive learning, a method that leverages positive and negative sample pairs to enhance model understanding. Positive pairs are selected via the Retrieval Tool introduced in Section 4.1.2. Specifically, for each anchor query, we retrieve its top-20 most

similar queries based on our hybrid similarity metric and treat them as positive samples. All remaining queries within the same batch serve as negative samples. This in-batch negative sampling strategy ensures scalability and efficiency, as it avoids the need to explicitly mine hard negatives from the full dataset while still providing sufficient contrastive signal.

After training, we apply the fine-tuned models to encode all queries from the clustering benchmarks. Pairwise distances between queries are computed using cosine similarity. We evaluate clustering quality using the BetaCV [36] metric, which measures intra-cluster compactness relative to inter-cluster distance. Specifically, a lower BetaCV score indicates better clustering performance, as it reflects that queries within the same cluster are more tightly grouped, while those from different clusters are more separated.

Table 5 presents the BetaCV scores across all models and benchmarks. Traditional rule-based methods such as Aouiche, Aligon, and Makiyama rely on handcrafted SQL features or syntax trees, and often fail to capture deep semantic similarities across structurally different but semantically related queries. PreQR improves upon this by leveraging a pretrained language model, but it lacks task-specific supervision and may underfit domain-specific patterns. Meanwhile, untuned BERT models benefit from general semantic priors but are still agnostic to SQL-specific structure and vocabulary. By contrast, our contrastively fine-tuned BERT models significantly outperform all baselines. Both versions of the BERT model have achieved the best performance on these benchmarks after training. These results validate the effectiveness of SQL-Factory in producing high-quality, semantically diverse SQL queries that can be used as training signals for representation learning.

## 6 CONCLUSION

In this paper, we propose SQL-Factory, a multi-agent framework designed to generate high-quality and large-scale SQL queries. By organizing the system into three collaborative teams for generation, expansion, and management, SQL-Factory achieves a balance between diversity, efficiency, and cost. The framework leverages powerful language models to explore novel structures, while lightweight local models efficiently scale promising patterns. Through real-time evaluation and schema-aware planning, SQL-Factory dynamically adjusts generation focus based on coverage gaps and schema complexity, ensuring that structurally rich tables receive sufficient attention. We demonstrate the effectiveness of SQL-Factory by applying it to multiple benchmarks, generating over 300,000 high-quality queries with minimal cost. Extensive experiments show that the generated queries significantly enhance downstream tasks such as Text-to-SQL and query clustering, confirming SQL-Factory as a practical solution for building large-scale, high-utility SQL corpora.

## ACKNOWLEDGMENTS

This work was supported in part by the NSFC under Grants No. (62302436, 62025206, U23A20296), Zhejiang Province’s “Lingyan” R&D Project under Grant No. 2024C01259, Yongjiang Talent Programme “Research on key technologies in Intelligent Database”, Ningbo Science and Technology Special Projects under Grant No. 2023Z212. Yuren Mao is the corresponding author of this paper.

## REFERENCES

- [1] Shadi Abdul Khalek and Sarfraz Khurshid. 2010. Automated SQL query generation for systematic testing of database engines. In *ASE*. 329–332.
- [2] Meta AI. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices>.
- [3] Julien Aligon, Matteo Golfarelli, Patrick Marcel, Stefano Rizzi, and Elisa Turricchia. 2014. Similarity measures for OLAP sessions. *Knowledge and information systems* 39, 2 (2014), 463–489.
- [4] Kamel Aouiche, Pierre-Emmanuel Jouve, and Jérôme Darmont. 2006. Clustering-based materialized view selection in data warehouses. In *ADBS*. 81–95.
- [5] Alexander Bianchi, Andrew Chai, Vincent Corvinelli, Parke Godfrey, Jarek Szlichta, and Calisto Zuzarte. 2024. Db2une: Tuning Under Pressure via Deep Learning. *Proc. VLDB Endow.* 17, 12 (2024), 3855–3868.
- [6] Jean-Flavien Bussotti, Enzo Veltri, Donatello Santoro, and Paolo Papotti. 2023. Generation of Training Examples for Tabular Natural Language Inference. In *SIGMOD*. 1–27.
- [7] Bikash Chandra, Bhupesh Chawda, Bipal Kar, K Reddy, Shetal Shah, and S Sudarshan. 2015. Data generation for testing and grading SQL queries. *Proc. VLDB Endow.* 24, 6 (2015), 731–755.
- [8] Weize Chen, Yusheng Su, Jingwei Zuo, et al. 2023. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors in agents. *arXiv preprint arXiv:2308.10848* (2023).
- [9] Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023. Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128* (2023).
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*. 4171–4186.
- [11] Bailu Ding, Surajit Chaudhuri, Johannes Gehrke, and Vivek Narasayya. 2021. DSB: A decision support benchmark for workload-driven and traditional database systems. *Proc. VLDB Endow.* 14, 13 (2021), 3376–3388.
- [12] Xuemei Dong, Chao Zhang, Yuhang Ge, Yuren Mao, Yunjun Gao, Jinshu Lin, Dongfang Lou, et al. 2023. C3: Zero-shot text-to-sql with chatgpt. *arXiv preprint arXiv:2307.07306* (2023).
- [13] Dan Friedman and Adji Bousso Dieng. 2022. The vendi score: A diversity evaluation metric for machine learning. *arXiv preprint arXiv:2210.02410* (2022).
- [14] 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 (2024), 1132–1145.
- [15] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 6894–6910.
- [16] Robert Giaquinto, Dejiao Zhang, Benjamin Kleiner, et al. 2023. Multitask pretraining with structured knowledge for text-to-SQL generation. In *ACL*. 11067–11083.
- [17] Daya Guo, Yibo Sun, Duyu Tang, et al. 2018. Question Generation from SQL Queries Improves Neural Semantic Parsing. In *EMNLP*. 1597–1607.
- [18] Daya Guo, Qihao Zhu, Dejian Yang, et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196* (2024).
- [19] Chenxu Hu, Jie Fu, Chenzhuang Du, Simian Luo, Junbo Zhao, and Hang Zhao. 2023. Chatdb: Augmenting llms with databases as their symbolic memory. *arXiv preprint arXiv:2306.03901* (2023).
- [20] Yi Huang, Fangyin Cheng, Fan Zhou, et al. 2024. ROMAS: A Role-Based Multi-Agent System for Database monitoring and Planning. *arXiv preprint arXiv:2412.13520* (2024).
- [21] Binyuan Hui, Jian Yang, Zeyu Cui, et al. 2024. Qwen2.5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [22] Oliver Kennedy, Jerry Ajay, Geoffrey Challen, and Lukasz Ziarek. 2016. Pocket data: The need for TPC-MOBILE. In *TPCTC*. 8–25.
- [23] Hideo Kobayashi, Wuwei Lan, Peng Shi, et al. 2025. You Only Read Once (YORO): Learning to Internalize Database Knowledge for Text-to-SQL. In *NAACL*. 1889–1901.
- [24] Igor Kotsyuba, Pavel Bezkorovaynyi, Julia Silko, and Alexey Shikov. 2021. Automating MySQL database complexity estimation based on quantitative metrics. In *INTERAGROMASH*. 378–386.
- [25] Gokhan Kul, Duc Thanh Anh Luong, Ting Xie, Varun Chandola, Oliver Kennedy, and Shambhu Upadhyaya. 2018. Similarity Metrics for SQL Query Clustering. *IEEE Transactions on Knowledge and Data Engineering* 30, 12 (2018), 2408–2420.
- [26] Viktor Leis, Bernhard Radke, Andrey Gubichev, et al. 2018. Query optimization through the looking glass, and what we found running the join order benchmark. *Proc. VLDB Endow.* 27, 5 (2018), 643–668.
- [27] Haoyang Li, Shang Wu, Xiaokang Zhang, et al. 2025. OmniSQL: Synthesizing High-quality Text-to-SQL Data at Scale.
- [28] Haoyang Li, Jing Zhang, Hanbing Liu, et al. 2024. CodeS: Towards Building Open-source Language Models for Text-to-SQL. In *SIGMOD*. 1–28.
- [29] Jinyang Li, Binyuan Hui, Ge Qu, et al. 2023. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls. *NeurIPS* (2023), 42330–42357.
- [30] Jiahui Li, Junhao Ye, Yuren Mao, Yunjun Gao, and Lu Chen. 2025. LOFTune: A Low-Overhead and Flexible Approach for Spark SQL Configuration Tuning. *IEEE Transactions on Knowledge and Data Engineering* 37, 6 (2025), 3528–3542.
- [31] Pengfei Li, Wenqing Wei, Rong Zhu, Bolin Ding, Jingren Zhou, and Hua Lu. 2023. ALECE: An Attention-based Learned Cardinality Estimator for SPJ Queries on Dynamic Workloads. *Proc. VLDB Endow.* 17, 2 (2023), 197–210.
- [32] Xinyue Liu, Xiangnan Kong, Lei Liu, and Kuorong Chiang. 2018. TreeGAN: syntax-aware sequence generation with generative adversarial networks. In *ICDM*. 1140–1145.
- [33] Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *ACL*. 142–150.
- [34] Vitor Hirota Makiyama, M Jordan Raddick, and Rafael DC Santos. 2015. Text Mining Applied to SQL Queries: A Case Study for the SDSS SkyServer. In *SIMBIG*. 66–72.
- [35] Ali Modarressi, Ayyoob Imani, Mohsen Fayyaz, and Hinrich Schütze. 2023. Ret-llm: Towards a general read-write memory for large language models. *arXiv preprint arXiv:2305.14322* (2023).
- [36] J Zaki Mohammed and Meira Wagner. 2014. Data mining and analysis: fundamental concepts and algorithms. *Cambridge University* (2014).
- [37] Raghunath Othayoth Nambiar and Meikel Poes. 2006. The making of TPC-DS. In *VLDB*, Vol. 6. 1049–1058.
- [38] Parimarjan Negi, Ziniu Wu, Andreas Kipf, et al. 2023. Robust query driven cardinality estimation under changing workloads. *Proc. VLDB Endow.* 16, 6 (2023), 1520–1533.
- [39] OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>.
- [40] Mile Pavlic, Marin Kaluza, and Neven Vreck. 2008. Database complexity measuring method. In *CECIS*. 577–583.
- [41] Mario Piattini, Coral Calero, and Marcela Genero. 2001. Table oriented metrics for relational databases. *Software Quality Journal* 9, 2 (2001), 79–97.
- [42] Meikel Poes and Chris Floyd. 2000. New TPC benchmarks for decision support and web commerce. *ACM SIGMOD Record* 29, 4 (2000), 64–71.
- [43] Mohammadreza Pourreza and Davood Rafiei. 2023. Din-sql: Decomposed in-context learning of text-to-sql with self-correction. *NeurIPS* (2023), 36339–36348.
- [44] Mohammadreza Pourreza, Davood Rafiei, Yuxi Feng, et al. 2024. SQL-Encoder: Improving NL2SQL In-Context Learning Through a Context-Aware Encoder.
- [45] Chen Qian, Wei Liu, Hongzhang Liu, et al. 2024. ChatDev: Communicative Agents for Software Development. In *ACL*. 15174–15186.
- [46] Andreas Seltenreich. 2016. SQLSmith. <https://github.com/anse1/sqlsmith>
- [47] Donald R Slutz. 1998. Massive stochastic testing of SQL. In *VLDB*, Vol. 98. 618–622.
- [48] Ruoxi Sun, Sercan Ö Arik, Alex Muzio, Lesly Miculich, Satya Gundabathula, Pengcheng Yin, Hanjun Dai, Hootan Nakhost, Rajarishi Sinha, Zifeng Wang, et al. 2023. Sql-palm: Improved large language model adaptation for text-to-sql (extended). *arXiv preprint arXiv:2306.00739* (2023).
- [49] Xiu Tang, Sai Wu, Mingli Song, Shanshan Ying, Feifei Li, and Gang Chen. 2022. PreQR: pre-training representation for SQL understanding. In *SIGMOD*. 204–216.
- [50] Xiangru Tang, Anni Zou, Zhuosheng Zhang, Ziming Li, Yilun Zhao, Xingyao Zhang, Arman Cohan, and Mark Gerstein. 2024. Medagents: Large language models as collaborators for zero-shot medical reasoning. In *ACL*. 599–621.
- [51] Dana Van Aken, Andrew Pavlo, Geoffrey J. Gordon, and Bohan Zhang. 2017. Automatic Database Management System Tuning Through Large-scale Machine Learning. In *SIGMOD*. 1009–1024.
- [52] Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, Linzheng Chai, Zhao Yan, Qian-Wen Zhang, Di Yin, Xing Sun, et al. 2025. Mac-sql: A multi-agent collaborative framework for text-to-sql. In *ICML*. 540–557.
- [53] Fang Wang, Xiao Yan, Man Lung Yiu, Shuai Li, Zunyao Mao, and Bo Tang. 2023. Speeding Up End-to-end Query Execution via Learning-based Progressive Cardinality Estimation. In *SIGMOD*. 1–25.
- [54] Lei Wang, Chen Ma, Xueyang Feng, et al. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science* 18, 6 (2024), 186345.
- [55] Taiyi Wang, Liang Liang, Guang Yang, Thomas Heinis, and Eiko Yoneki. 2025. A New Paradigm in Tuning Learned Indexes: A Reinforcement Learning Enhanced Approach. *arXiv preprint arXiv:2502.05001* (2025).
- [56] Kun Wu, Lijie Wang, Zhenghua Li, et al. 2021. Data Augmentation with Hierarchical SQL-to-Question Generation for Cross-domain Text-to-SQL Parsing. In *EMNLP*. 8974–8983.
- [57] Zhiheng Xi, Wenxiang Chen, Xin Guo, et al. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences* 68, 2 (2025), 121101.
- [58] Wenxuan Xie, Gaochen Wu, and Bowen Zhou. 2024. Mag-sql: Multi-agent generative approach with soft schema linking and iterative sub-sql refinement for text-to-sql. *arXiv preprint arXiv:2408.07930* (2024).
- [59] Wenyi Xu, Yuren Mao, Xiaolu Zhang, et al. 2025. DAgent: A Relational Database-Driven Data Analysis Report Generation Agent. *arXiv preprint arXiv:2503.13269* (2025).

- [60] Tao Yu, Rui Zhang, Kai Yang, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task. In *EMNLP*. 3911–3921.
- [61] Tao Yu, Zhaonian Zou, and Hao Xiong. 2024. Can Uncertainty Quantification Enable Better Learning-based Index Tuning? *arXiv preprint arXiv:2410.17748* (2024).
- [62] Chao Zhang, Yuren Mao, Yijiang Fan, et al. 2024. FinSQL: model-agnostic LLMs-based text-to-SQL framework for financial analysis. In *SIGMOD*. 93–105.
- [63] Ji Zhang, Yu Liu, Ke Zhou, et al. 2019. An End-to-End Automatic Cloud Database Tuning System Using Deep Reinforcement Learning. In *SIGMOD*. 415–432.
- [64] Lixi Zhang, Chengliang Chai, Xuanhe Zhou, and Guoliang Li. 2022. Learned-SQLGen: Constraint-aware SQL Generation using Reinforcement Learning. In *SIGMOD*. 945–958.
- [65] Yi Zhang, Jan Deriu, George Katsogiannis-Meimarakis, et al. 2023. ScienceBenchmark: A Complex Real-World Benchmark for Evaluating Natural Language to SQL Systems. *Proceedings of the VLDB Endowment* 17, 4 (2023), 685–698.
- [66] Yaowei Zheng, Richong Zhang, Junhao Zhang, et al. 2024. LlamaFactory: Unified Efficient Fine-Tuning of 100+ Language Models. In *ACL*. 400–410.
- [67] Zhiling Zheng, Oufan Zhang, Ha L Nguyen, et al. 2023. Chatgpt research group for optimizing the crystallinity of mofs and cofs. *ACS Central Science* 9, 11 (2023), 2161–2170.
- [68] Xuanhe Zhou, Guoliang Li, Zhaoyan Sun, et al. 2024. D-Bot: Database Diagnosis System using Large Language Models. *Proc. VLDB Endow.* 17, 10 (2024), 2514–2527.