



SQL-Exchange: Transforming SQL Queries Across Domains

Mohammadreza Daviran
University of Alberta
Edmonton, AB, Canada
daviran@ualberta.ca

Brian Lin
University of Alberta
Edmonton, AB, Canada
shihhsu1@ualberta.ca

Davood Rafiei
University of Alberta
Edmonton, AB, Canada
drafie@ualberta.ca

ABSTRACT

We introduce SQL-Exchange, a framework for mapping SQL queries across different database schemas by preserving the source query structure while adapting domain-specific elements to align with the target schema. We investigate the conditions under which such mappings are feasible and beneficial, and examine their impact on enhancing the in-context learning performance of text-to-SQL systems as a downstream task. Our comprehensive evaluation across multiple model families and benchmark datasets—assessing structural alignment with source queries, execution validity on target databases, and semantic correctness—demonstrates that SQL-Exchange is effective across a wide range of schemas and query types. Our results further show that both in-context prompting with mapped queries and fine-tuning on mapped data consistently yield higher text-to-SQL performance than using examples drawn directly from the source schema.

PVLDB Reference Format:

Mohammadreza Daviran, Brian Lin, and Davood Rafiei. SQL-Exchange: Transforming SQL Queries Across Domains. PVLDB, 19(6): 1291 - 1304, 2026.

doi:10.14778/3797919.3797935

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/mmdrez4/SQL-Exchange>.

1 INTRODUCTION

What do *toxicology* and *Formula 1* have in common, and where do *superheroes* intersect with *student clubs*? These are all databases in the BIRD benchmark [18], and despite the differences in the semantics of their tables and columns, the structure of their queries is similar, and in many cases, identical. Consider the following example queries across three different databases:

- (1) What is the total number of clients in the Vsetin district?
- (2) What is the total number of atoms in molecules with a label of '+'?
- (3) What is the total number of races held at the Canadian Grand Prix circuit?

As shown in Figure 1, all three queries share the same SQL skeleton despite the differences in their query semantics and schema links. The question studied in this paper is if queries expressed on a source database can be mapped to equivalent queries on a target

database. Here, we use the term ‘equivalence’ loosely to refer to queries that are structurally identical, independent of schema links and constants.

Although the source and target domains may differ semantically, structurally aligned queries are valuable in several practical scenarios. Many real databases, especially newly created or under-documented ones, lack sufficient NL–SQL pairs for in-context learning or fine-tuning, yet structurally relevant examples are known to significantly improve performance; for example, Pourreza et al. [35] reports that using in-context examples tailored for the target domain improves SQL generation performance by 12% compared to using cross-domain examples. SQL-Exchange enables the automatic creation of schema-aligned SQL–NL pairs directly for the target schema, supporting prompting, data augmentation, and model adaptation in low-resource settings. Structural mappings are also useful for bootstrapping query repositories: in the SDSS project [1], a repository of queries was constructed for a new database to support the search and reuse of similar queries [7, 14]. Additionally, educational and assessment settings [13, 27] benefit from structurally parallel queries, where varying the domain while preserving logical structure reduces plagiarism and helps maintain consistent difficulty across assignments.

Existing work on data migration—an industry valued at USD 8.2 billion in 2021 and projected to reach USD 33.58 billion by 2030 [30]—has traditionally focused on transforming data and queries across different database engines and architectures (e.g., SQL-to-NoSQL, SQL Server to SQLite, and SQL to Hive [40]). In contrast, text-to-SQL research predominantly focuses on translating natural language into SQL under a single, fixed schema [4, 17, 33]. Parallel efforts explore other query translation tasks, such as SQL-to-English [25], and SQL-to-NoSQL [2], aimed at improving accessibility or enabling cross-platform execution. However, most prior work focuses on syntactic translation or query generation rather than preserving the structural logic of SQL queries while adapting them to new database schemas. To the best of our knowledge, no existing work systematically addresses schema-level SQL query translation—that is, mapping queries to databases with different schemas while preserving their structural integrity. We address this gap with a novel LLM-based framework that maps SQL queries across schemas, maintaining the original query structure and enabling seamless transfer of logic between databases.

There are intriguing questions and challenges in mapping queries across domains. First, queries that are relevant in one domain may not be meaningful in another due to differences in the number and semantics of tables and columns, as well as the presence or absence of foreign key relationships and constraints. Second, as query complexity increases, they may involve multiple joins, nested subqueries, aggregations, and filtering conditions, all of which introduce dependencies that must be correctly adapted to the target

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 19, No. 6 ISSN 2150-8097.
doi:10.14778/3797919.3797935

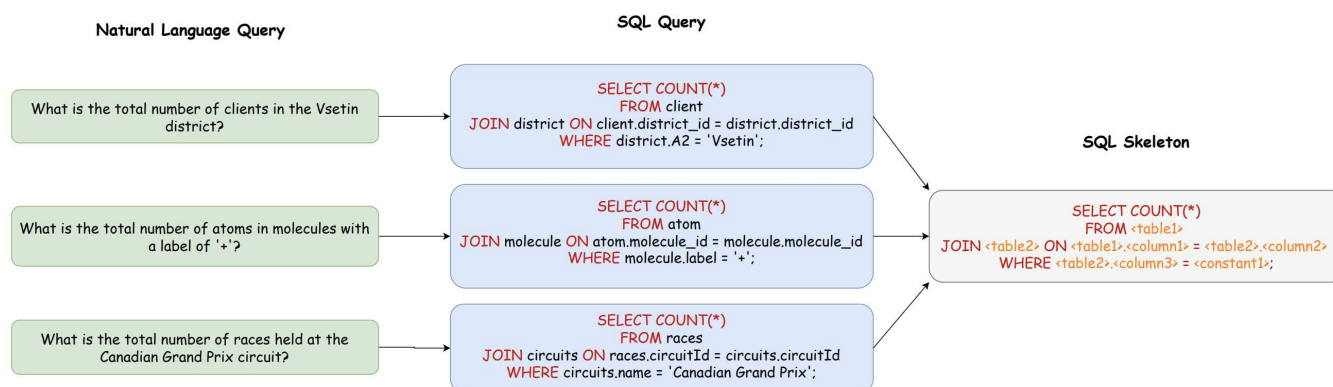


Figure 1: Illustration of three natural language queries, their corresponding SQL translations, and their shared SQL skeleton, demonstrating structural similarity across different database schemas.

schema, making it more difficult to transfer the query logic to a different domain. Furthermore, the relationships between tables and columns in different schemas may be structurally different, requiring transformations that go beyond simple column substitutions. The impact of these factors on query portability is neither well-studied nor well-understood.

In this paper, we (1) introduce a novel approach for mapping SQL queries across different domains, while preserving their logical structure, and (2) conduct a comprehensive empirical investigation into its behavior and limitations. This work aims to address three key questions: (1) Can SQL queries be accurately mapped across different database schemas while preserving their structural form and logic? (2) Can LLMs serve as reliable tools for performing such mappings? and (3) Do mapped queries effectively improve the performance in downstream tasks such as text-to-SQL generation?

To answer these questions, we undertake an extensive, multi-dimensional evaluation that characterizes the structural fidelity, execution validity, and natural-language meaningfulness of the mapped queries across BIRD and SPIDER. Our analysis reveals that SQL-NL pairs generated through SQL-Exchange exhibit strong structural alignment (over 82%) and high NL quality (over 95%) using Gemini-1.5-flash, with GPT-4o-mini showing similarly strong trends. These findings are supported by small-scale manual evaluations and larger-scale assessments using LLM-as-a-judge.

Our contributions can be summarized as follows: (1) We introduce and analyze a framework for cross-domain SQL query mapping that preserves the structural logic of a source query while adapting its schema-dependent components to a target database. (2) Through extensive evaluation on BIRD and SPIDER using multiple LLM families, we show that SQL-Exchange produces queries that remain structurally faithful, executable, and semantically meaningful. We also identify the main patterns of structural adaptation and the key factors that affect mapping quality. (3) We construct a corpus of over 100,000 schema-aligned SQL-NL pairs and demonstrate their downstream utility for *text-to-SQL*: SQL-Exchange examples consistently outperform unmapped and zero-shot baselines across retrieval methods, and fine-tuning on mapped examples yields substantially larger gains than fine-tuning on source-only data, underscoring the method’s value in low-supervision settings.

2 RELATED WORK

Synthetic Data Generation with LLMs. LLMs have been extensively used for synthetic data generation across various natural language processing tasks, including classification, data augmentation, and code generation [3, 21, 29, 43]. These efforts typically aim to expand the volume and diversity of training data by generating paraphrases or novel examples under controlled conditions. A particular relevant area of application is question generation over structured knowledge bases and narrative texts [20, 22, 44]. For instance, Li and Zhang [20] introduce a planning-first strategy that guides LLMs in generating questions with greater control over both structure and content. Most of these methods focus on generating isolated, free-form questions with minimal structural constraints—significantly different from the more complex task of generating SQL-NL pairs, which requires maintaining consistency between natural language and formal query logic.

Synthetic SQL Generation and Cross-Dialect Translation. Early work on synthetic data generation for text-to-SQL employed template-based SQL synthesis and copy-based Seq2Seq models to translate SQL into natural language [5, 8, 39, 42]. More recent approaches leverage LLMs to generate SQL-NL pairs grounded in real database schemas [26, 36]. Large-scale pipelines such as OmniSQL [16], SQL-Factory [19], and SQLForge [6] further expand training corpora through multi-agent generation, AST-level composition, and broad schema exploration.

These systems, however, are designed for data expansion: they generate new SQL queries—often across thousands of synthetic or isolated schemas—to support large-scale pretraining. In contrast, SQL-Exchange does not aim to enlarge training sets or synthesize new schemas. Instead, it transforms an existing SQL query into another schema while preserving its structural scaffold. This enables domain-adapted one-shot or few-shot in-context examples, a capability not supported by synthetic generation pipelines whose outputs are not tied to the logical structure of a specific source query. Whereas synthetic generators optimize for coverage and diversity, SQL-Exchange preserves cross-domain structural alignment.

A parallel line of work addresses SQL dialect or engine translation [23, 31, 47], using rule-based methods, traditional models, or LLM-assisted pipeline to improve query portability. These approaches typically assume identical schemas across dialects or target structurally simple SQL queries, limiting their applicability to schema-adapted transformations.

Schema Mapping and Data Migration. Data migration—the process of moving data between formats, systems, or platforms—has a long history in the database and systems communities [11, 15, 24, 28]. When data spans multiple databases with differing schema representations, a central challenge is managing semantic heterogeneity across database schemas [11, 12].

Schema conversion is often employed in SQL-to-NoSQL migrations, particularly to improve join performance by reorganizing data so that join-relevant attributes reside in the same document or collection [46]. In NoSQL systems, schema evolution may occur lazily, driven by application interactions and without requiring downtime or disrupting active workloads [37]. Query transformations are also frequently necessary; for example, translating SQL to Hive often requires modifying queries when certain SQL constructs are not supported by the target system [40].

Unlike traditional data migration techniques, which typically preserve the original data and query semantics despite changes in representation or syntax, cross-domain query transformation—as considered in this work—allows both data values and query intent to shift. In our setting, the source query serves as a structural guide rather than a strict equivalence, enabling adaptation to target schemas with differing semantics or domain-specific conventions.

3 METHODOLOGY

We aim to translate a natural language query and its corresponding SQL from a source schema to an equivalent query pair on a target schema while preserving the original logical structure.

3.1 Problem Formulation

Let q_s be a natural language question and sq_s its corresponding SQL query over a source schema s . Given a target schema t , our goal is to generate a natural language question q_t and an SQL query sq_t over t such that:

- q_t is a meaningful and well-formed question over t ;
- sq_t correctly expresses the semantics of q_t within t ;
- sq_s and sq_t share the same logical structure, abstracting away from table names, column names, and constants.

3.2 Limitations of Zero-Shot Translation

Large Language Models (LLMs) have demonstrated impressive capabilities in translation tasks ranging from natural language to formal representations (e.g., text-to-SQL [4, 33]) to cross-format conversions (e.g., SQL to NoSQL [2] and tabular data representations [9, 32]). To assess their suitability for schema-to-schema SQL translation, we began with a zero-shot setting in which the model received only a high-level task description and schema metadata.

Given (q_s, sq_s, s, t) , the LLM was asked to generate (q_t, sq_t) directly. This setup revealed two critical challenges:

Structural Drift. The model frequently failed to preserve the structural backbone of the source SQL. Structural alignment—the degree to which the high-level logic (e.g., joins, aggregations, filters) is retained—was often below 50%, and in some benchmarks, below 15%. A common issue was the omission of essential JOIN clauses, flattening multi-table logic into oversimplified single-table queries.

Schema Leakage. The model also showed a tendency to copy table names, column names, and constants from the source schema into the target query, without adaptation. For instance, GPT-4o-mini reused 15% of source column names and over 70% of literal values in one dataset, often leading to syntactically incorrect or semantically irrelevant queries.

BIRD: appstore → debit_card_specializing	
Source SQL:	SELECT AVG(Price) FROM playstore WHERE Genres = 'Dating'
Target SQL:	SELECT AVG(T1.Price) FROM products AS T1 WHERE T1.Description = 'Dating'
BIRD: soccer_2016 → toxicology	
Source SQL:	SELECT T2.Outcome_Type FROM Match AS T1 INNER JOIN Outcome AS T2 ON T1.Outcome_type = T2.Outcome_Id WHERE T1.Match_Id = '392195'
Target SQL:	SELECT T1.label FROM molecule AS T1 WHERE T1.molecule_id = '392195'

Figure 2: Examples of structural drift and schema leakage in zero-shot mapping using Gemini. In the first case, schema elements (in red)—such as column names and literals—are inappropriately copied from the source into the target query (e.g., no price column in target). In the second case, the target query omits a necessary JOIN clause from the source, leading to a loss of structural fidelity.

These challenges, illustrated in Figure 2, motivated the development of a more structured prompting strategy, as discussed next.

3.3 Structured Prompting with Chain-of-Thought Reasoning

To address the weaknesses of zero-shot prompting, we adopt a structured one-shot prompting approach in which the model receives: (1) a high-level task description; (2) one illustrative example demonstrating a correct source-to-target query mapping; and (3) the source and target schema definitions together with the source query.

This format capitalizes on the benefits of in-context learning while remaining within LLM context limits.

Additionally, inspired by chain-of-thought (COT) prompting [41], we incorporate an intermediate reasoning step. Rather than immediately generating sq_t , the model first produces a brief natural language explanation—a *thought process*—that outlines how it interprets the query logic and how it plans to map it to the target schema. This step enhances interpretability and guides structurally faithful SQL generation. The complete structured prompt, including instructions and the required inputs, is shown in the box below:

Structured Prompt with Reasoning
<p>Input: A natural language question (q_s) and SQL query (sq_s) over a source schema (s); a target schema (t); sample data from the target schema (d_t); and a demonstration example (e) from another schema pair.</p> <p>Output: A JSON object containing a mapped natural language question (q_t) and SQL query (sq_t) over the target schema (t).</p> <hr/> <p>## System prompt: You are an expert in areas of database design and SQL queries and your job is to swap tables and columns from a query to generate a new query, and new question based on the target schema.</p> <p>## Instructions:</p> <ul style="list-style-type: none"> - The output must be in a valid JSON format. - For each Source query, there must be a corresponding output in the output array. - First, create <code>tables_columns_replacement</code> by replacing all table names, column names, and constant values in the source query with placeholders: "table", "column", "constant_value". - Then, perform the following steps and give your thought process (in not more than 5 sentences) for each step: <ol style="list-style-type: none"> 1: Generate a new query from <code>tables_columns_replacement</code>. <ol style="list-style-type: none"> 1.1: Use Target schema to replace table and column names that make sense in terms of query. 1.2: For constant values in <code>tables_columns_replacement</code>: <ul style="list-style-type: none"> • Use meaningful values from the target schema sample data. • Do not reuse constant values from the "source_query" • Ensure that numerical constant values differ from those in the "source_query". 2: Generate a new question based on the query that you just generated. <p>## Example: {e}</p> <p>## Generate the query for the following query:</p> <p># Source schema: s # Target schema: t # Target sample data: d_t # Source query: {q_s, sq_s}</p> <p># Output:</p>

The overall pipeline includes two key mechanisms described below: query template abstraction and constant substitution.

3.4 Template-Guided Query Transfer

One major failure in zero-shot settings is the model’s difficulty in disentangling schema-independent logic from schema-specific artifacts, often leading to poor structural alignment and inappropriate reuse of tables or columns from the source schema. To address this, we introduce a template-guided abstraction mechanism that helps the model disentangle logical structure from surface-level artifacts.

The process begins by converting the source SQL query sq_s into a schema-agnostic **query template**, where all table names, column names, and constants are replaced with generic placeholders such as `table`, `column`, and `value`. Crucially, this abstraction preserves the structural backbone of the query—including joins, aggregations, and filters—while removing database-specific identifiers. For example, a query that joins two tables on a foreign key and applies a filtering condition becomes a high-level template with the same logical form but anonymized components. This abstraction encourages the model to focus on generalizable relational patterns rather than memorized schema tokens.

Once the template is created, the model grounds the template by instantiating it with appropriate elements from the target schema t . This step involves selecting valid table and column names and substituting meaningful constants, resulting in a complete SQL query sq_t that aligns with both the logic of the source query and the structure of the target database. This two-phase approach promotes compositional reasoning, reduces schema leakage, and yields structurally faithful translations.

Empirical results confirm the effectiveness of this approach. On the BIRD benchmark, for example, using Gemini-1.5-Flash, structural alignment improved from just 13.1% in the zero-shot setting to 66.6% with template-guided prompting. Similar gains were observed across other models and benchmarks, demonstrating the robustness of this strategy.

To illustrate, consider this example from a schema on postal and demographic data (with structural elements highlighted in blue):

Source Schema

NL Question: Provide the names of bad aliases in the city of Aguadilla.
SQL: `SELECT T1.bad_alias FROM avoid AS T1 INNER JOIN zip_data AS T2 ON T1.zip_code = T2.zip_code WHERE T2.city = 'Aguadilla'`

Its abstract template becomes:

Schema-Agnostic Template

`SELECT T1.column FROM table1 AS T1 INNER JOIN table2 AS T2 ON T1.column2 = T2.column2 WHERE T2.column3 = constant_value`

Within the target chemical-compounds schema, this maps to:

Target Schema

NL Question: What are the bond types in molecule 'TR028'?
SQL: `SELECT T1.bond_type FROM bond AS T1 INNER JOIN molecule AS T2 ON T1.molecule_id = T2.molecule_id WHERE T2.molecule_id = 'TR028'`

This illustrates how structure is retained while adapting the query to an entirely different domain.

3.5 Substituting Constant Values

Substituting constant values poses a significant challenge in query mapping. Unlike table and column names, which can often be aligned through structural reasoning over the target schema, constants—such as categorical labels (e.g., district names, molecule types) or numerical values—are highly dataset-specific and rarely transferable across domains.

To address this, we augment the prompt with a small set of sample rows from the target database. This gives the model access to valid constant values grounded in the target schema. Additionally, we explicitly instruct the LLM to replace each constant placeholder with a semantically appropriate value that exists in the target database, rather than copying constants from the source query. Without such guidance or access to target-side data, LLMs tend to default to reusing source query constants, leading to semantically invalid outputs that may be syntactically correct but inconsistent with the target schema. This was especially evident in our zero-shot experiments, where constant reuse rates reached as high as 72% on BIRD using GPT-4o-mini, indicating that most target queries failed to generate schema-appropriate values. In contrast, SQL-Exchange reduced constant reuse by more than 30 percentage points, producing values that matched the target schema distribution and yielding higher execution validity and question quality.

4 EVALUATION

We evaluate the performance of SQL-Exchange across multiple model families, datasets, and evaluation metrics.

4.1 Experimental Setup

Datasets We evaluate SQL-Exchange on two standard text-to-SQL benchmarks: *BIRD* [18] and *SPIDER* [45]. For each benchmark, we select one source database from the training set and one target database from the development set, applying the same instructions and prompt structure across all models and database pairs. BIRD includes 69 training and 11 development databases, while SPIDER includes 146 training and 20 development databases. We sample up to 20 queries per source database; for GPT-4o-mini on SPIDER, we reduce this to 10 for cost reasons. If a database has fewer than the desired number of queries, all available queries are included.

To verify that our sampling does not bias the evaluation toward simpler SQL structures, we examined the *hardness* of the sampled subsets using the official Spider rule-based criteria (easy, medium, hard, extra-hard), defined by structural markers such as joins, nesting, grouping, and set operations. For BIRD, the sampled hardness distribution closely matches the full dataset, with differences below 0.6% in every category; a chi-square test further confirms no statistical difference ($\chi^2 = 0.50$, $p = 0.918$). For SPIDER, the differences remain modest (below 2.8% per category). Although the chi-square test detects a shift ($\chi^2 = 16.14$, $p = 0.001$), the absolute deviations are small, and the sample still contains substantial numbers of hard and extra-hard queries, preserving the overall complexity profile of the full dataset.

LLMs for Mapping. For the query mapping task, we employed two LLMs from distinct model families: GPT-4o-mini, Gemini-1.5-flash [38], and the open-weight Llama-3.3-70B. The first two models were selected for their fast response times, cost-efficiency, and

long context windows. Both offer 128k-token context windows, enabling us to batch 10 to 20 queries per prompt—alongside long schema descriptions and sample inputs—thereby reducing the number of LLM calls and minimizing token redundancy. The combined prompt length often exceeded the 16k-token context window of models such as GPT-3.5-turbo. In contrast, both GPT-4o-mini and Gemini-1.5-flash handled such inputs reliably with minimal failures. We additionally included Llama-3.3-70B-Instruct to evaluate SQL-Exchange under an open model. Throughout the paper, we refer to model–dataset combinations using shorthand notations such as BIRD–Gemini, BIRD–GPT, SPIDER–Gemini, and BIRD–Llama.

LLM for Evaluation. To assess semantic quality, we use Gemini-2.0-flash as an LLM-based evaluator due to its superior reasoning performance. This model is validated against human annotators for assessing the correctness and quality of generated queries.

Model Access and Configuration. GPT-4o-mini was accessed through the OpenAI API,¹ Gemini models via Google’s developer API,² and Llama-3.3-70B through the Nebius API.³ All models were queried with temperature = 0.0, top_p = 1.0, and top_k = 0 whenever supported.

Batching and Prompting. LLM prompts included the task instruction, full source and target schema descriptions, target-side sample rows, and a set of up to 10-20 source queries. Prompts for query mapping were batched wherever possible to reduce API calls.

4.2 Metrics

We report results across two dimensions: (i) mapping success and fidelity, and (ii) semantic plausibility of the generated queries.

Mapping Accuracy. We first measure how reliably the model can produce mapped queries under the SQL-Exchange using three metrics:

- **Generation Success:** Percentage of queries for which the LLM generates a non-empty, parseable response instead of refusing or skipping the mapping.
- **Structural Alignment:** Whether the mapped query preserves the structural skeleton of the source SQL—including keywords, control blocks (e.g., SELECT, JOIN, GROUP BY, sub-queries), and overall logic. Table and column names, constants, and optional AS aliasing are ignored, and basic comparison operators (e.g., =, <, >) are normalized.
- **Execution Validity:** Whether the generated SQL executes successfully on the target schema using SQLite.

Semantic Quality. We evaluate whether the generated queries are semantically valid via:

- **NL Meaningfulness:** Whether the generated natural language question is clear, specific, and meaningful in the context of the target database schema.
- **SQL–NL Alignment:** Whether the SQL query faithfully implements the question’s intent.

¹<https://platform.openai.com/docs/models/gpt-4o-mini>

²<https://ai.google.dev/models/gemini>

³<https://tokenfactory.nebius.com/>

Table 1: Evaluation results on the BIRD and SPIDER benchmarks (for Gemini-1.5-flash and GPT-4o-mini). Semantic quality metrics—NL Meaningfulness and SQL–NL Alignment—are based on an LLM-as-a-judge valuation using Gemini-2.0-flash. Manual evaluation results on a smaller 144-query subset are reported in Table 2.

Category	Metric	BIRD Dataset		SPIDER Dataset	
		Gemini-1.5-flash	GPT-4o-mini	Gemini-1.5-flash	GPT-4o-mini
Mapping Accuracy	Generation Success	99.47%	100.00%	98.9%	100.00%
	Structural Alignment	66.59%	80.47%	82.21%	86.8%
	Execution Validity	89.43%	68.33%	93.5%	86.29%
Semantic Quality	NL Meaningfulness	95.45%	80.03%	96.41%	91.78%
	SQL–NL Alignment	82.75%	54.57%	90.02%	76.13%

4.3 Mapping Accuracy Analysis

Table 1 summarizes aggregate results on BIRD and SPIDER, reporting mean performance for each benchmark across mapping accuracy and semantic quality. We next analyze the quality of the SQL–NL pairs produced by SQL-Exchange.

4.3.1 Generation Success. A mapping query was produced in nearly all cases, with failures—defined as explicit null outputs—occurring in under 1% of instances. This behaviour varied across LLMs: GPT-4o-mini consistently generated a mapping for every input, whereas Gemini-1.5-flash occasionally returned null values, particularly for complex queries or when structural incompatibilities existed between the source and target schemas.

For example, one failure case involved a source query in BIRD with an invalid HAVING clause combining MAX and MIN without proper aggregation logic:

Failure Case 1 (Invalid logic)

```
SELECT ... GROUP BY ... HAVING MAX(ibu) AND MIN(ibu) LIMIT 2
```

Here, the model correctly identified the logical flaw and refused to produce a misleading translation.

In another case, the source query identified the employee with the most inspections in March 2016:

Failure Case 2 (Schema mismatch)

```
SELECT ... FROM (
  SELECT employee_id, COUNT(...) FROM ...
  WHERE strftime('%Y-%m', ...) = '2016-03'
  GROUP BY employee_id
  ORDER BY COUNT(...) DESC LIMIT 1
) AS T2 INNER JOIN ... ON ...
```

This query combines temporal filtering, aggregation, ordering, and a join to retrieve personnel metadata. The model declined to generate a corresponding query for the toxicology schema, noting that the source logic was “difficult to directly translate to the simpler structure of the target schema.” Importantly, the toxicology schema contains no temporal attributes or event logs (as highlighted in red), making the transformation fundamentally infeasible. Rather than

hallucinating a misleading mapping, the model conservatively returned a null output. This behavior reflects SQL-Exchange’s robustness in recognizing when faithful transformation is not possible.

4.3.2 Structural Alignment. In 66% to 87% of cases, the mapped queries accurately preserved the structural skeleton of the source SQL queries while adapting to the target schema. Perfect structural preservation is not always feasible due to differences between source and target schemas, which may require modifying joins, groupings, or aggregation functions to ensure semantic alignment.

To better understand how LLMs adapt query structures during mapping, we categorize token-level edits into ten semantic buckets based on SQL keywords: JOIN (join clauses), COND (filters and logical conditions), DISTINCT (keyword), FUNC_OTHER (scalar, formatting, and null-handling functions), AGG_FUNC (aggregate functions), GROUPING (grouping and filtering), ORDER (ordering and row limits), SUBQUERY (nested queries and set operations), PUNCT_SELECT (commas in SELECT clause indicating changes in the number of output columns), and OTHER (residual edits). Each change is labeled as a deletion, insertion, or mutation, and may fall into multiple categories. For instance, editing a fragment like JOIN ... WHERE ... IN (SELECT ...) may trigger JOIN, COND, and SUBQUERY simultaneously.

Figure 3 presents the distribution of these edits for Gemini-1.5-flash and GPT-4o-mini on BIRD and SPIDER. Deletions dominate, followed by insertions and fewer mutations. This trend suggests that when LLMs modify query structure, they often simplify queries by removing joins, filters, or modifiers that do not directly align with the target schema.

We observe that the majority of structural edits occur in the JOIN and COND buckets. These include deletions of unnecessary joins or filters and mutations that generalize conditions (e.g., transforming ‘=’ to ‘LIKE’). While the model frequently simplifies joins, it also introduces conditions or filters to adapt queries to the target schemas. Gemini tends to perform more insertions in FUNC_OTHER, AGG_FUNC, and SUBQUERY, such as adding STRFTIME(), LENGTH(), or COUNT() to improve compatibility with target schema semantics. Similar patterns are observed in GPT-4o-mini mappings, though Gemini tends to perform more structural substitutions involving functions and aggregations.

SPIDER queries exhibit a stronger tendency toward deletion, particularly in JOIN and COND categories. This is likely due to the




Figure 3: Frequency of structural changes, broken down to SQL constructs, for the two benchmarks (BIRD, SPIDER) and two LLMs (GEMINI-1.5-FLASH, GPT-4O-MINI). Each bar is stacked by **deletions**, **insertions**, and **mutations**. Since a single query may involve multiple types of edits across different clauses, bucket totals may exceed the number of modified queries. This visualization reveals which parts of the SQL skeleton are most frequently altered during schema adaptation. Note that the charts are not on the same scale, and the dataset used for the SPIDER-Gemini setting is larger than that for SPIDER-GPT.

higher degree of normalization and the presence of auxiliary tables in SPIDER schemas. These characteristics can make certain joins redundant or overly complex in the mapped queries.

In terms of overall change patterns, deletions account for 55.2% of structural edits, followed by insertions (32.2%) and mutations (12.6%). These statistics reinforce the observation that LLMs simplify query structure when faced with schema mismatches, even though the overall skeleton is often preserved.

To complement the aggregate statistics, Figure 4 presents representative examples of structural edits made during query mapping. These illustrate SQL-Exchange’s ability to adapt queries to diverse schemas: (A) introducing temporal functions to bridge column-type mismatches, such as adapting a year comparison (`built_year < 1850`) to a timestamp field using `strftime('%Y', created) < '2019'`; and (B) removing joins when the target schema provides a direct column that captures the same semantic condition, allowing the query to be expressed more simply. These edits reflect the model’s schema-sensitive and structure-aware behavior.

4.3.3 Execution Validity. As shown in Table 1, execution success rates for mapped queries range from 68.3% to 93.5%, across models and datasets, with Gemini-1.5-flash performing best. This indicates that, in most cases, the mapped queries are both syntactically correct and semantically well-aligned with the target schema.

To better understand the failure cases, we categorized all SQL error messages into six classes. Most errors are **column reference errors** (90.5–93.4%), typically involving incorrect table–column associations, nonexistent columns, or missing quotes for identifiers containing spaces. A smaller portion consists of **syntax errors** (1.9–6.3%)—such as missing commas, stray parentheses, or unescaped identifiers—which appear more frequently in GPT-4o-mini outputs. Other errors include **table reference mistakes** (0.3–3.2%) from incorrect or hallucinated table names; **ambiguous column names** (0.7–2.5%), which occur when join-related columns lack disambiguation; and **misuses of aggregate functions** (0.2–2.5%), such as invalid or nested SUM, MAX, or COUNT. The remaining **rare errors** ($\leq 1.2\%$) involve malformed set operators,

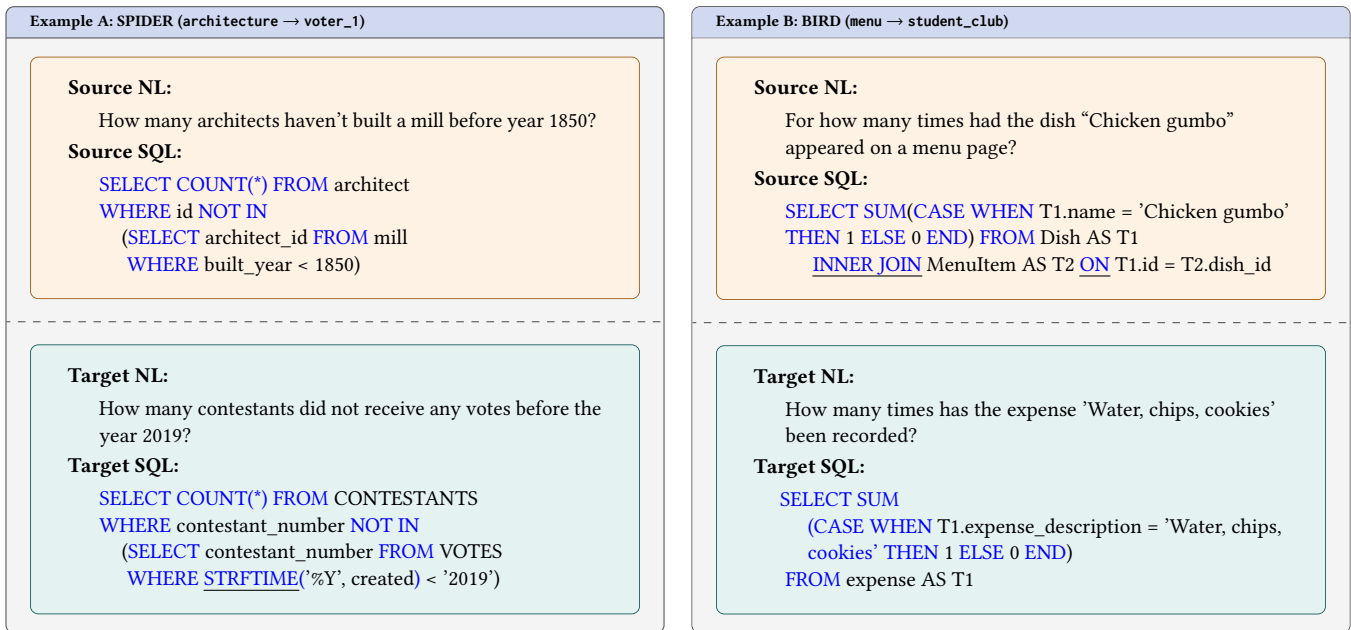


Figure 4: Representative examples of structural edits during query mapping by SQL-Exchange. (A) Introduction of the STRFTIME function to enable temporal filtering on string-based timestamps; (B) Removal of a join when the target schema supports a direct column-based formulation of the counting logic. Examples are drawn from the SPIDER and BIRD training sets and mapped to target schemas from their respective development sets.

incomplete SELECT clauses, improperly nested subqueries, or timeouts caused by complex joins or missing indexes. This breakdown shows that even structurally correct queries can fail due to subtle schema mismatches, suggesting that future work may benefit from schema-level validation or improved quoting mechanisms.

4.3.4 Structure vs. Utility. While structural alignment is a key goal of our method, Table 1 shows that preserving structure alone does not guarantee better performance. GPT-4o-mini exhibits higher structural alignment than Gemini-1.5-flash on both benchmarks, yet performs significantly worse in execution validity and semantic alignment, especially on BIRD. This suggests that strict structural preservation can hurt semantic consistency or produce non-executable SQL. In contrast, Gemini-1.5-flash makes small structural adjustments when needed, producing more executable queries that better align with the natural language intent. These findings indicate that effective cross-schema mapping requires balancing structural fidelity with semantic adaptability.

4.3.5 Open-Weight Model Evaluation. While Table 1 reports full-benchmark results only for Gemini-1.5-flash and GPT-4o-mini, we also conducted a brief open-weight model evaluation using Llama-3.3-70B. Since BIRD is substantially more complex than SPIDER, this evaluation was limited to the same BIRD subset used for the Gemini and GPT models. On this subset, Llama achieved competitive performance across all metrics (98.51% generation success, 85.23% structural alignment, 70.93% execution validity, 85.55% NL meaningfulness, and 63.54% SQL-NL alignment). These results demonstrate that SQL-Exchange is effective under modern open-weight models as well.

4.4 Semantic Quality of Generated Queries

To evaluate the semantic quality of the mapped queries, we assess whether the generated natural language questions are meaningful within the context of the target schema, and whether the corresponding SQL queries are logically consistent with the intent of the generated question. This evaluation is performed both manually and through an LLM-based review.

4.4.1 Manual Evaluation. We conducted a manual evaluation on a curated subset of 144 SQL-NL pairs that passed execution testing. To ensure diverse coverage, we selected three target databases that exhibited low execution success rates or weak structural alignment in our earlier analyses: *california_schools* and *european_football_2* (from BIRD), and *poker_player* (from SPIDER). For each, we randomly sampled eight source schemas and selected three structurally varied queries per pair—categorized as simple (≤ 5 elements), moderate (6–9), or complex (≥ 10), based on the number of tables and columns. This produced 72 mappings, evaluated under two LLMs (GPT-4o-mini and Gemini-1.5-flash), yielding 144 total examples.

Each of the three authors independently judged whether (1) the NL question was meaningful for the target schema and (2) the SQL query matched that question’s intent. Only cases with agreement from at least two annotators were included—resulting in 138 judgments for NL quality and 134 for SQL-NL alignment.

As shown in Table 2, across all settings, the vast majority of generated questions (100% for Gemini-1.5-flash and 88–91% for GPT-4o-mini) were judged meaningful. SQL-NL alignment was high for Gemini-1.5-flash (91–96%) and moderate for GPT-4o-mini (56–57%), highlighting a notable gap in fidelity across models.

Table 2: Manual evaluation of 144 mapped queries, assessing NL meaningfulness and SQL-NL alignment. Annotator agreement was reached at 138 queries for NL meaningfulness and 134 for SQL-NL alignment. Trends align with LLM-based results in Table 1.

Benchmark-Model	Meaningful NL	Aligned SQL
BIRD-Gemini	100.0%	91.1%
BIRD-GPT	88.6%	57.1%
SPIDER-Gemini	100.0%	95.8%
SPIDER-GPT	91.3%	56.5%

4.4.2 Validating LLM-as-a-Judge Against Human Annotations. To assess the potential of LLMs as scalable evaluators, we used Gemini-2.0-flash to score the same 144 query pairs evaluated manually. As with the manual evaluation, we report results only for the subset of queries where at least two out of three annotators agreed: 138 for NL meaningfulness and 134 for SQL-NL alignment. Gemini-2.0-flash achieved **94.9%** agreement on question meaningfulness and **85.07%** on SQL-NL alignment, closely matching human annotations. Most disagreements were either borderline “maybe” cases, due to subjective interpretation. In the full disagreements (e.g., one labeled as “yes” and the other as “no”), the LLM’s explanation was often more consistent and grounded than the human rationale, suggesting that LLM-based evaluation can serve as a reliable alternative to manual review.

4.4.3 Scaling Up with LLM-as-a-Judge. After Gemini-2.0-flash was validated against human judgment, we used it to evaluate the full set of mapped queries. As shown in Table 1, most generated NL questions were judged meaningful (95–96% for Gemini-1.5-flash and 80–91% for GPT-4o-mini). SQL-NL alignment was similarly strong for Gemini-1.5-flash (82–90%), and more variable for GPT-4o-mini (54–76%), reflecting a consistent performance gap across models and datasets. Importantly, this evaluation was conducted on the complete set of generated queries, regardless of whether they passed structural alignment or execution validity checks. This allows us to assess the semantic plausibility even in structurally or syntactically imperfect outputs, which may explain the slightly lower scores compared to the manually curated subset.

4.5 Ablation: SQL-Exchange vs. Zero-Shot

To assess the contribution of each component in our methodology, we compare SQL-EXCHANGE with a zero-shot prompting baseline. Both approaches use the same LLMs and schema information, but the baseline omits key elements of our framework: structured Chain-of-Thought, structural query templates, demonstration examples, and target schema sample rows. This comparison isolates the impact of structural abstraction and schema grounding on mapping quality.

We evaluate both approaches using the core metrics introduced in Table 1, and additionally introduce a new metric, **Result Yield**, which measures whether the generated query returns non-empty results on the target database. This metric complements execution validity by assessing whether the query is not only syntactically correct but also semantically grounded in the target schema’s data.

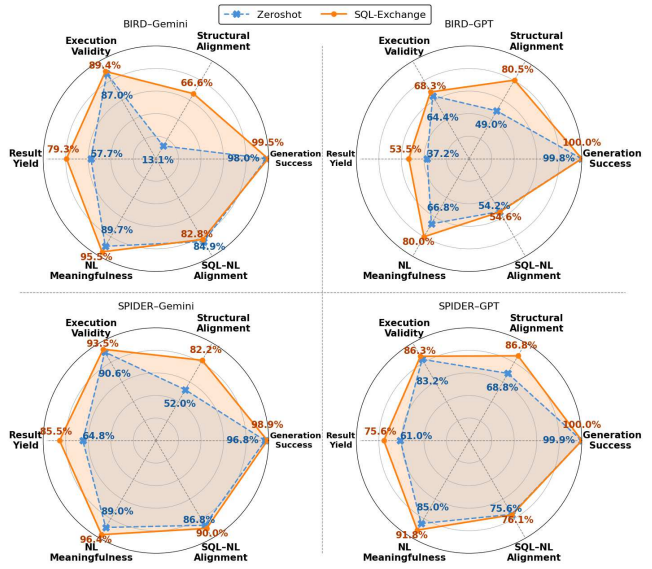


Figure 5: Comparison of SQL-Exchange and Zero-Shot prompting across semantic and structural metrics.

Figure 5 shows four radar plots comparing SQL-Exchange and zero-shot prompting across six evaluation metrics, one plot for each model–dataset setting. SQL-Exchange demonstrates consistent gains in both structural and semantic quality: structural alignment improves by +18% to +53.53%, result yield increases by +14.7% to +21.7%, and meaningful questions rates improve by +5.3% to +13.3%. SQL-Exchange also consistently outperforms zero-shot on both *Generation Success* and *Execution Validity* across all settings, underscoring the effectiveness of SQL-Exchange’s abstraction-driven templates and schema-grounded constant substitution mechanisms.

Performance on *NL–SQL alignment* is more mixed. Zero-shot has a slight advantage in the BIRD-Gemini setting, but SQL-Exchange performs better in the remaining three. Even when alignment scores appear similar, zero-shot queries often simplify the structure by omitting essential components such as JOIN clauses or conditional logic. Although the resulting queries may remain executable or logically aligned in intent, they diverge from the original compositional semantics. In contrast, SQL-Exchange preserves structural scaffolding far more faithfully.

We further analyze SQL-Exchange’s behavior relative to the zero-shot baseline along three dimensions: (1) structural simplification through deletions, (2) preservation of relational joins, and (3) avoidance of undesirable reuse from the source schema.

Structural simplification via deletion. On average, zero-shot prompting exhibited significantly more deletions than the mappings produced by SQL-Exchange—ranging from approximately 1.5× to 4.5× depending on the model and benchmark. The most frequently deleted components were JOIN clauses and conditional filters, which are essential for preserving multi-table logic and row-level constraints. This indicates a strong simplification trend, where zero-shot prompting omits critical structural elements rather than faithfully adapting them.

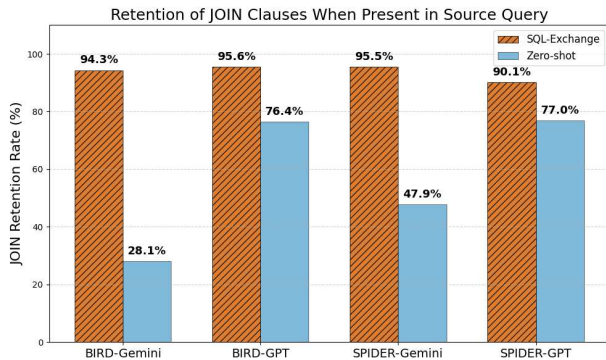


Figure 6: Join retention rates: the proportion of mapped queries that preserve a JOIN present in the source. SQL-Exchange retains joins consistently (>90%), while zero-shot often drops them, especially with Gemini.

Preservation of relational joins. As shown in Figure 6, zero-shot prompting often fails to retain JOIN clauses in the mapped output when the source query originally included one, resulting in structurally oversimplified outputs. On BIRD-Gemini and SPIDER-Gemini, fewer than 30% and 50% of such queries, respectively, contain any join clause after mapping. In contrast, SQL-Exchange retains joins in over 90% of cases across all settings, suggesting that schema abstraction in SQL-Exchange enables more faithful multi-table reasoning when adapting query structure across domains.

Avoidance of undesirable schema reuse. Figure 7 shows that SQL-Exchange substantially reduces the reuse of constants, tables, and columns from the source schema—an issue that plagues zero-shot prompting. Table reuse decreases by over 90% relative to the zero-shot baseline in all settings, falling below 0.5%. Column reuse drops by 62–77% relative to zero-shot, demonstrating consistent suppression of schema-copying artifacts. This confirms that SQL-Exchange’s schema-aware design not only preserves structural logic more reliably, but also minimizes unintended copying of irrelevant source-specific elements.

5 DOWNSTREAM TASK PERFORMANCE

Beyond structural and semantic correctness, we also assess whether the mapped queries are useful as in-context examples or fine-tuning data in downstream text-to-SQL tasks. The key question is whether these schema-aligned queries improve SQL generation accuracy in one-shot/few-shot prompting or when fine-tuning local models.

5.1 Experimental Setup

We evaluate three instruction-tuned open-source models: Llama-3.2-3B-Instruct, Qwen2.5-Coder-3B-Instruct, and Qwen2.5-Coder-7B-Instruct [10]. In addition, for supervised fine-tuning experiments, we use Mistral-7B-Instruct-v0.3 and CodeLlama-7b-Instruct. Throughout the paper, these models are referred to as LLaMA-3B, Qwen-3B, Qwen-7B, Mistral-7B, and CodeLlama-7B, respectively. We also include results from the proprietary models GPT-4o-mini and GPT-4o.

Inference. Open-source models were run locally with Hugging Face Transformers on a single NVIDIA A100 (40GB) using float16. Inference used greedy decoding (temperature = 0.0, top_p = 1.0) and instruction-style formatting. Fine-tuning was performed on Google Colab with an 80GB A100. GPT models were accessed via the OpenAI API with the same greedy decoding setup.

In-Context Example Selection. We evaluated multiple strategies for selecting in-context examples, including embedding-based retrieval, masked-question selection, DAIL [4], and SQL-Encoder⁴. SQL-Encoder is a 1.3B-parameter model trained to estimate structural similarity between SQL queries using only their natural language questions [35]. For all selection methods, candidates were retrieved under two standard settings: (i) from source-side queries (unmapped) and (ii) from mapped outputs produced by SQL-Exchange. To obtain an approximate upper bound on achievable performance, we additionally used SQL-Encoder to retrieve examples from the target development sets as well. Retrieval was performed independently for each model-dataset setting and filtered by semantic and execution validity where applicable.

5.2 Metrics

We measure downstream performance using **execution accuracy**, the standard text-to-SQL metric (percentage of generated queries matching the gold execution result). To analyze quality control during example selection, we also report **filtering ablations**, removing semantic and execution filters to isolate their contribution. Semantic filtering removes NL-SQL mismatches, while execution filtering excludes non-executable queries.

5.3 Prompting Strategies and Results

We evaluate a broad set of prompting strategies, grouped by how in-context examples are selected: **Zero-shot**, which generates SQL without examples; **Embedding-based (1-shot)**, which retrieves the closest example using Sentence Transformers⁵; and **Masked Question (1-shot)**, which retrieves examples using a masked-question representation of the NL input. We also include **DAIL (1-shot)** [4], which selects examples via Query Similarity Selection; to do this, we use either a preliminary model (GPT-4o-mini, [†]) or the gold SQL as an oracle ([◇]) to identify the best example from each candidate pool. Finally, we use **SQL-Encoder (1-shot / 3-shot)**, which retrieves structurally similar examples using NL-based structural similarity, and we apply it in both 1-shot and 3-shot settings as well as in an oracle mode using target dev-set pairs.

Table 3 reports execution accuracy on BIRD and SPIDER across all models and selection methods. As seen throughout the table, unmapped examples frequently underperform even zero-shot, especially for smaller models, whereas SQL-Exchange consistently provides stronger demonstrations that better transfer across schemas.

Across nearly all retrieval methods and models, SQL-Exchange substantially outperforms Unmapped examples. On SPIDER, gains frequently exceeding +5 EX and reaching as high as +19.05 EX for LLaMA-3B under SQL-Encoder (3-shot). On BIRD, improvements similarly reach up to +11.93 EX. For GPT-4o, improvements are

⁴<https://huggingface.co/MrezaPRZ/sql-encoder>

⁵We use the all-mpnet-base-v2 model for embedding-based retrieval.

Reuse Rate and Relative Reduction for Harmful Behaviors




Figure 7: Reduction in schema-specific reuse across settings. We compare reuse rates for constants, tables, and columns between zero-shot prompting and SQL-Exchange. Each subplot shows actual reuse percentages across four schema settings, with SQL-Exchange in solid bars and zero-shot in transparent bars. Above each SQL-Exchange bar, we annotate the relative reduction (%) compared to the zero-shot baseline. SQL-Exchange dramatically lowers the rate of constant, table, and column reuse.

Table 3: Execution accuracy (EX) on BIRD and SPIDER development sets across LLMs and example-selection methods. Results are reported as BIRD / SPIDER. SQL-Exchange examples consistently improve over zero-shot and unmapped baselines, and SQL-Encoder with oracle retrieval provides an approximate upper bound.

Selection Method	Prompting Strategy	LLaMA-3B BIRD/SPIDER	Qwen-3B BIRD/SPIDER	Qwen-7B BIRD/SPIDER	GPT-4o-mini BIRD/SPIDER	GPT-4o BIRD/SPIDER
<i>No In-Context Examples</i>	Zero-shot	26.53 / 43.33	35.98 / 53.48	49.80 / 71.08	50.46 / 66.73	56.39 / 69.34
<i>Embedding-based</i>	Unmapped (1-shot)	22.88 / 40.52	32.66 / 52.80	47.26 / 62.57	50.39 / 67.21	58.67 / 70.70
	SQL-Exchange (1-shot)	32.40 / 52.61	39.37 / 60.15	50.52 / 67.22	52.48 / 72.15	58.21 / 73.98
<i>Masked Question</i>	Unmapped (1-shot)	23.60 / 42.94	33.57 / 53.38	46.94 / 64.41	49.54 / 69.92	58.60 / 71.66
	SQL-Exchange (1-shot)	31.68 / 52.13	39.70 / 59.28	50.39 / 66.34	52.35 / 69.83	59.58 / 73.98
<i>DAIL [4]</i>	Unmapped [†] (1-shot)	23.92 / 42.75	34.68 / 56.58	48.63 / 65.57	50.33 / 69.63	59.19 / 71.76
	SQL-Exchange [†] (1-shot)	31.42 / 50.97	38.40 / 59.19	49.35 / 65.96	51.24 / 67.50	58.02 / 71.86
	Unmapped [◇] (1-shot)	24.05 / 46.23	35.53 / 58.90	47.00 / 65.76	51.04 / 72.73	60.69 / 75.43
	SQL-Exchange [◇] (1-shot)	33.64 / 53.97	42.11 / 63.83	49.48 / 69.34	52.28 / 73.79	59.91 / 77.66
<i>SQL-Encoder</i>	Unmapped (1-shot)	25.03 / 43.42	31.94 / 60.93	46.94 / 70.12	48.37 / 68.86	58.74 / 70.60
	SQL-Exchange (1-shot)	35.33 / 57.64	41.46 / 70.70	51.89 / 73.60	51.50 / 69.73	58.67 / 68.28
	Unmapped (3-shot)	20.73 / 37.62	34.29 / 53.97	46.87 / 71.66	49.15 / 69.44	60.50 / 72.63
	SQL-Exchange (3-shot)	32.66 / 56.67	43.09 / 68.67	50.85 / 73.89	53.19 / 73.40	60.17 / 78.05
	Oracle (1-shot)	39.05 / 76.02	48.89 / 84.82	58.21 / 90.14	57.11 / 84.04	63.10 / 80.56
	Oracle (3-shot)	34.94 / 74.66	49.15 / 87.52	59.65 / 90.04	58.60 / 86.17	65.78 / 86.26

[†] DAIL performed using GPT-4o-mini for query similarity selection.

[◇] DAIL performed using the gold SQL query as an oracle for query similarity selection.

smaller in a few settings, not because SQL-Exchange is ineffective, but because demonstrations were generated using substantially weaker models (Gemini-1.5-flash and GPT-4o-mini). We expect larger gains when SQL-Exchange examples are produced using a model comparable in strength to GPT-4o. Given that most SQL-Exchange vs. Unmapped differences exceed the natural variance of SPIDER ($\approx 1-1.5$ EX) and BIRD ($\approx 1.5-2$ EX), paired significance tests confirm that the majority of SQL-Exchange improvements are statistically significant (typically $p < 0.05$, often $p < 0.01$), with only very small deltas (< 2 EX) falling in the nonsignificant range.

In the SQL-Encoder setting, both 1-shot and 3-shot methods generally outperform their unmapped counterparts when using SQL-Exchange. Although we apply 3-shot prompting only within SQL-Encoder, results show that adding more schema-aligned examples helps larger models (e.g., +5.4% for GPT-4o on SPIDER), while smaller models sometimes see mild declines due to added noise.

Overall, SQL-Exchange provides effective and robust in-context examples across model scales, retrieval methods, and prompting strategies, without requiring task-specific fine-tuning or large synthetic corpora.

Table 4: Ablation study showing the effect of removing execution or semantic filtering on accuracy for both BIRD and SPIDER development sets in one-shot and few-shot settings. Semantic filtering contributes most to downstream performance across both prompting strategies.

Model	Configuration	BIRD	SPIDER
LLaMA-3B	SQL-Exchange (1-shot)	35.33%	57.64%
	w/o execution filter	34.03%	57.83%
	w/o semantic filter	33.96%	55.61%
	SQL-Exchange (3-shot)	32.66%	56.67%
	w/o execution filter	32.20%	56.87%
Qwen-3B	w/o semantic filter	32.33%	53.87%
	SQL-Exchange (1-shot)	41.46%	70.7%
	w/o execution filter	40.48%	70.6%
	w/o semantic filter	40.74%	69.34%
	SQL-Exchange (3-shot)	43.09%	68.67%
Qwen-7B	w/o execution filter	41.26%	69.25%
	w/o semantic filter	42.05%	67.12%
	SQL-Exchange (1-shot)	51.89%	73.6%
	w/o execution filter	51.17%	73.4%
	w/o semantic filter	50.78%	72.24%
GPT-4o-mini	SQL-Exchange (3-shot)	50.85%	73.89%
	w/o execution filter	50.39%	73.6%
	w/o semantic filter	50.39%	73.69%
	SQL-Exchange (1-shot)	51.5%	69.73%
	w/o execution filter	51.56%	70.21%
GPT-4o-mini	w/o semantic filter	52.74%	70.21%
	SQL-Exchange (3-shot)	53.19%	73.4%
	w/o execution filter	53.39%	74.27%
	w/o semantic filter	52.93%	74.47%

Filtering Ablations. Table 4 examines the effect of removing execution or semantic filtering when selecting mapped examples. We restrict our ablation study to the SQL-Encoder setting, which supports both 1-shot and 3-shot prompting and offers the most complete retrieval pipeline, making it a representative case for isolating filter effects. Extending ablations to all methods and models would require substantially more LLM calls, which is impractical in time and cost. For the same reason, GPT-4o is not included.

Across models, removing semantic filtering causes the largest accuracy drops—especially for smaller LMs (e.g., -1.4% on BIRD for LLaMA-3B, -1.36% on SPIDER for Qwen-3B). This shows that keeping the NL question aligned with the mapped SQL is crucial. Execution filtering has a smaller impact and can even help slightly when removed for stronger models (e.g., $+0.5\%$ for GPT-4o-mini), indicating that larger LMs can tolerate minor SQL imperfections. Smaller models, however, are more sensitive to both types of errors, underscoring the value of filtering for lower-capacity settings.

Understanding the Oracle-Mapped Gap. While our method consistently outperforms zero-shot and unmapped one-shot baselines, it still falls short of Oracle performance. This gap arises because Oracle examples, drawn from the development set, often

closely mirrors the target query logic—sharing joins, filters, ordering, and aggregation. In contrast, SQL-Exchange mappings are constrained by the availability and structure of source queries, which may lack strong logical or schema-level alignment.

Mapped queries sometimes retain redundant joins (e.g., auxiliary tables) or introduce irrelevant filters, misleading the LLM. Oracle examples also employ more precise constructs (e.g., IS NOT NULL, DISTINCT) and exact aggregation logic that better match the NL question. Lastly, Oracle queries often reflect more natural phrasing and keyword overlap with the test NL question, further helping the model follow the correct reasoning path.

5.4 Finetuning

To further assess the utility of SQL-Exchange, we fine-tune downstream models on two of the most challenging BIRD development databases—*california_schools* and *toxicology*—where baseline performance is lowest. Following the setup of [34], we fine-tune Mistral-7B on two datasets: (1) our 1.3K sampled source-database queries from the BIRD training set, and (2) 1.3K SQL-Exchange-generated mappings created by transferring those queries into the target schema. On *california_schools*, zero-shot Mistral-7B achieves **6.74%** execution accuracy; fine-tuning on source-only data reduces accuracy to **5.62%**, whereas fine-tuning on SQL-Exchange-generated data boosts it to **14.61%**. A similar trend appears on *toxicology*, with accuracy improving from **8.28%** (zero-shot) to **10.34%** (source-only) and **22.07%** with SQL-Exchange-generated data.

Similar behavior is observed with CodeLlama-7B. On *california_schools*, accuracy moves from **3.37%** zero-shot to **1.12%** with source-only fine-tuning, but rises to **7.87%** when fine-tuned on SQL-Exchange mappings. On *toxicology*, performance increases from **8.97%** to **15.17%** with source-only fine-tuning and to **20.00%** with SQL-Exchange data. These results demonstrate that SQL-Exchange provides substantially more effective supervision on difficult target schemas than source-only data.

6 CONCLUSION

We introduced a novel schema-aware method for mapping SQL queries across database schemas using LLMs, a problem not directly addressed in prior work. Our approach preserves the structural skeleton of the source query while adapting schema-specific elements and constants through one-shot prompting, template-based abstraction, and sample-guided substitution. Our evaluations on BIRD and SPIDER show that most of the mapped queries are structurally aligned with the source, executable on the target, and semantically consistent with the generated natural language. Furthermore, when used as in-context examples in downstream text-to-SQL tasks, these mappings improve execution accuracy over zero-shot and unmapped baselines, and fine-tuning on SQL-Exchange-generated mappings provides substantial additional gains on challenging databases compared to fine-tuning on unmapped queries or relying on zero-shot prompting.

ACKNOWLEDGMENTS

This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). Computational resources were provided by the Digital Research Alliance of Canada.

REFERENCES

- [1] A.P. Sloan Foundation. [n.d.]. SDSS sky server query search. <https://skyserver.sdss.org/dr18/SearchTools/sql>. Accessed: 2024-08-09.
- [2] Wu-Chun Chung, Hung-Pin Lin, Shih-Chang Chen, Mon-Fong Jiang, and Yeh-Ching Chung. 2014. JackHare: a framework for SQL to NoSQL translation using MapReduce. *Automated Software Engineering* 21 (2014), 489–508.
- [3] Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Yihan Cao, Zihao Wu, Lin Zhao, Shaochen Xu, Fang Zeng, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, Hongmin Cai, Lichao Sun, Quanzheng Li, Dinggang Shen, Tianming Liu, and Xiang Li. 2025. AugGPT: Leveraging ChatGPT for Text Data Augmentation. *IEEE Transactions on Big Data* 11, 3 (2025), 907–918. <https://doi.org/10.1109/TBDATA.2025.3536934>
- [4] 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. *Proceedings of the VLDB Endowment* 17, 5 (2024), 1132–1145.
- [5] Daya Guo, Yibo Sun, Duyu Tang, Nan Duan, Jian Yin, Hong Chi, James Cao, Peng Chen, and Ming Zhou. 2018. Question Generation from SQL Queries Improves Neural Semantic Parsing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 1597–1607. <https://doi.org/10.18653/v1/D18-1188>
- [6] Yu Guo, Dong Jin, Shenghao Ye, Shuangwu Chen, Jian Yang, and Xiaobin Tan. 2025. SQLForge: Synthesizing Reliable and Diverse Data to Enhance Text-to-SQL Reasoning in LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar (Eds.). Association for Computational Linguistics, Vienna, Austria, 8441–8452. <https://doi.org/10.18653/v1/2025.findings-acl443>
- [7] Bill Howe, Garret Cole, Nodira Khoussainova, and Leilani Battle. 2011. Automatic example queries for ad hoc databases. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data (Athens, Greece) (SIGMOD '11)*. Association for Computing Machinery, New York, NY, USA, 1319–1322. <https://doi.org/10.1145/1989323.1989487>
- [8] Yiqun Hu, Yiyun Zhao, Jiarong Jiang, Wuwei Lan, Henghui Zhu, Anuj Chauhan, Alexander Hanbo Li, Lin Pan, Jun Wang, Chung-Wei Hang, Sheng Zhang, Jiang Guo, Mingwen Dong, Joseph Lilien, Patrick Ng, Zhiguo Wang, Vittorio Castelli, and Bing Xiang. 2023. Importance of Synthesizing High-quality Data for Text-to-SQL Parsing. In *Findings of the Association for Computational Linguistics: ACL 2023*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 1327–1343. <https://doi.org/10.18653/v1/2023.findings-acl.86>
- [9] Zezhou Huang, Jia Guo, and Eugene Wu. 2024. Transform table to database using large language models. *Proceedings of the VLDB Endowment*. ISSN 2150 (2024), 8097.
- [10] Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, et al. 2024. Qwen2.5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [11] Richard Hull. 1984. Relative information capacity of simple relational database schemata. In *Proceedings of the 3rd ACM SIGACT-SIGMOD Symposium on Principles of Database Systems (Waterloo, Ontario, Canada) (PODS '84)*. Association for Computing Machinery, New York, NY, USA, 97–109. <https://doi.org/10.1145/588011.588027>
- [12] Richard Hull. 1997. Managing semantic heterogeneity in databases: a theoretical perspective. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (Tucson, Arizona, USA) (PODS '97)*. Association for Computing Machinery, New York, NY, USA, 51–61. <https://doi.org/10.1145/263661.263668>
- [13] Breanna Jury, Angela Lorusso, Juho Leimonen, Paul Denny, and Andrew Luxton-Reilly. 2024. Evaluating LLM-generated Worked Examples in an Introductory Programming Course. In *Proceedings of the 26th Australasian Computing Education Conference (Sydney, NSW, Australia) (ACE '24)*. Association for Computing Machinery, New York, NY, USA, 77–86. <https://doi.org/10.1145/3636243.3636252>
- [14] Nodira Khoussainova, YongChul Kwon, Wei-Ting Liao, Magdalena Balazinska, Wolfgang Gatterbauer, and Dan Suciu. 2011. Session-based browsing for more effective query reuse. In *Proceedings of the 23rd International Conference on Scientific and Statistical Database Management (Portland, OR) (SSDBM '11)*. Springer-Verlag, Berlin, Heidelberg, 583–585.
- [15] Samir Khuller, Yoo-Ah Kim, and Yung-Chun (Justin) Wan. 2003. Algorithms for data migration with cloning. In *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (San Diego, California) (PODS '03)*. Association for Computing Machinery, New York, NY, USA, 27–36. <https://doi.org/10.1145/773153.773156>
- [16] Haoyang Li, Shang Wu, Xiaokang Zhang, Xinmei Huang, Jing Zhang, Fuxin Jiang, Shuai Wang, Tiejing Zhang, Jianjun Chen, Rui Shi, Hong Chen, and Cuiping Li. 2025. OmniSQL: Synthesizing High-Quality Text-to-SQL Data at Scale. *Proc. VLDB Endow.* 18, 11 (July 2025), 4695–4709. <https://doi.org/10.14778/3749646.3749723>
- [17] Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. 2023. RESDSL: decoupling schema linking and skeleton parsing for text-to-SQL. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence (AAAI'23/IAAI'23/EAAI'23)*. AAAI Press, Article 1466, 9 pages. <https://doi.org/10.1609/aaai.v37i11.26535>
- [18] Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin C.C. Chang, Fei Huang, Reynold Cheng, and Yongbin Li. 2023. Can LLM already serve as a database interface? a big bench for large-scale database grounded text-to-SQLs. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 1835, 28 pages.
- [19] Jiahui Li, Tongwang Wu, Yuren Mao, Yunjun Gao, Yajie Feng, and Huaizhong Liu. 2025. SQL-Factory: A Multi-Agent Framework for High-Quality and Large-Scale SQL Generation. *arXiv preprint arXiv:2504.14837* (2025).
- [20] Kunze Li and Yu Zhang. 2024. Planning First, Question Second: An LLM-Guided Method for Controllable Question Generation. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 4715–4729. <https://doi.org/10.18653/v1/2024.findings-acl.280>
- [21] Zhuoyan Li, Hangxiao Zhu, Zhuoran Lu, and Ming Yin. 2023. Synthetic Data Generation with Large Language Models for Text Classification: Potential and Limitations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 10443–10461. <https://doi.org/10.18653/v1/2023.emnlp-main.647>
- [22] Yuanyuan Liang, Jianing Wang, Hanlun Zhu, Lei Wang, Weining Qian, and Yunshi Lan. 2023. Prompting Large Language Models with Chain-of-Thought for Few-Shot Knowledge Base Question Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 4329–4343. <https://doi.org/10.18653/v1/2023.emnlp-main.263>
- [23] Zhisheng Lin, Yifu Liu, Zhiling Luo, Jinyang Gao, and Yu Li. 2024. MoMQ: Mixture-of-Experts Enhances Multi-Dialect Query Generation across Relational and Non-Relational Databases. *arXiv preprint arXiv:2410.18406* (2024).
- [24] Chenyang Lu, Guillermo A. Alvarez, and John Wilkes. 2002. Aqueduct: Online Data Migration with Performance Guarantees. In *Proceedings of the 1st USENIX Conference on File and Storage Technologies (Monterey, CA) (FAST '02)*. USENIX Association, USA, 21–es.
- [25] Wo-Shun Luk and Steve Kloster. 1986. ELFS: English language from SQL. *ACM Transactions on Database Systems (TODS)* 11, 4 (1986), 447–472.
- [26] Alisia Lupidi, Carlos Gemmel, Nicola Cancedda, Jane Dwivedi-Yu, Jason Weston, Jakob Foerster, Roberta Raileanu, and Maria Lomeli. 2024. Source2synth: Synthetic data generation and curation grounded in real data sources. *arXiv preprint arXiv:2409.08239* (2024).
- [27] Niklas Meißner, Sandro Speth, Julian Kieslinger, and Steffen Becker. 2024. Evalquiz—llm-based automated generation of self-assessment quizzes in software engineering education. In *Software Engineering im Unterricht der Hochschulen 2024*. Gesellschaft für Informatik eV, 53–64.
- [28] Renée J Miller, Yannis E Ioannidis, and Raghu Ramakrishnan. 1993. The use of information capacity in schema integration and translation. In *VLDB*, Vol. 93, 120–133.
- [29] Mihai Nadas, Laura Diosan, and Andreea Tomescu. 2025. Synthetic Data Generation Using Large Language Models: Advances in Text and Code. *arXiv preprint arXiv:2503.14023* (2025).
- [30] Next Move Strategy Consulting. 2023. Data migration market. <https://www.nextmsc.com/report/data-migration-market>. Accessed: 2024-08-16.
- [31] Amadou Latyr Ngom and Tim Kraska. 2024. Mallet: SQL Dialect Translation with LLM Rule Generation. In *Proceedings of the Seventh International Workshop on Exploiting Artificial Intelligence Techniques for Data Management (Santiago, AA, Chile) (aiDM '24)*. Association for Computing Machinery, New York, NY, USA, Article 3, 5 pages. <https://doi.org/10.1145/3663742.3663973>
- [32] Arash Dargahi Nobari and Davood Rafiei. 2025. TabulaX: Leveraging Large Language Models for Multi-Class Table Transformations. *Proc. VLDB Endow.* 18, 11 (July 2025), 3826–3839. <https://doi.org/10.14778/3749646.3749657>
- [33] Mohammadreza Pourreza and Davood Rafiei. 2024. Din-sql: Decomposed in-context learning of text-to-sql with self-correction. *Advances in Neural Information Processing Systems* 36 (2024).
- [34] Mohammadreza Pourreza and Davood Rafiei. 2024. DTS-SQL: Decomposed Text-to-SQL with Small Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 8212–8220. <https://doi.org/10.18653/v1/2024.findings-emnlp.481>
- [35] Mohammadreza Pourreza, Davood Rafiei, Yuxi Feng, Raymond Li, Zhenan Fan, and Weiwei Zhang. 2024. SQL-Encoder: Improving NL2SQL In-Context Learning Through a Context-Aware Encoder. *arXiv preprint arXiv:2403.16204* (2024).

- [36] Mohammadreza Pourreza, Ruoxi Sun, Hailong Li, Lesly Miculicich, Tomas Pfister, and Sercan O Arik. 2024. Sql-gen: Bridging the dialect gap for text-to-sql via synthetic data and model merging. *arXiv preprint arXiv:2408.12733* (2024).
- [37] Karla Saur, Tudor Dumitras, and Michael Hicks. 2016. Evolving NoSQL Databases without Downtime. In *2016 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 166–176. <https://doi.org/10.1109/ICSME.2016.47>
- [38] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530* (2024).
- [39] Bailin Wang, Wenpeng Yin, Xi Victoria Lin, and Caiming Xiong. 2021. Learning to Synthesize Data for Semantic Parsing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, Online, 2760–2766. <https://doi.org/10.18653/v1/2021.naacl-main.220>
- [40] Yue Wang, Yingzhong Xu, Yue Liu, Jian Chen, and Songlin Hu. 2015. QMapper for Smart Grid: Migrating SQL-based Application to Hive. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD '15)*. Association for Computing Machinery, New York, NY, USA, 647–658. <https://doi.org/10.1145/2723372.2742792>
- [41] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '22)*. Curran Associates Inc., Red Hook, NY, USA, Article 1800, 14 pages.
- [42] Kun Wu, Lijie Wang, Zhenghua Li, Ao Zhang, Xinyan Xiao, Hua Wu, Min Zhang, and Haifeng Wang. 2021. Data Augmentation with Hierarchical SQL-to-Question Generation for Cross-domain Text-to-SQL Parsing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 8974–8983. <https://doi.org/10.18653/v1/2021.emnlp-main.707>
- [43] Kang Min Yoo, Dongju Park, Jaewook Kang, Sang-Woo Lee, and Woomyoung Park. 2021. GPT3Mix: Leveraging Large-scale Language Models for Text Augmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 2225–2239. <https://doi.org/10.18653/v1/2021.findings-emnlp.192>
- [44] Hokeun Yoon and JinYeong Bak. 2023. Diversity Enhanced Narrative Question Generation for Storybooks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 465–482. <https://doi.org/10.18653/v1/2023.emnlp-main.31>
- [45] Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (Eds.). Association for Computational Linguistics, Brussels, Belgium, 3911–3921. <https://doi.org/10.18653/v1/D18-1425>
- [46] Gansen Zhao, Qiaoying Lin, Libo Li, and Zijing Li. 2014. Schema Conversion Model of SQL Database to NoSQL. In *Proceedings of the 2014 Ninth International Conference on P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC '14)*. IEEE Computer Society, USA, 355–362. <https://doi.org/10.1109/3PGCIC.2014.137>
- [47] Ran Zmigrod, Salwa Alami, and Xiaomo Liu. 2024. Translating between SQL Dialects for Cloud Migration. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice (Lisbon, Portugal) (ICSE-SEIP '24)*. Association for Computing Machinery, New York, NY, USA, 189–191. <https://doi.org/10.1145/3639477.3639727>