

000 SPOTIT^Q: EVALUATING TEXT-TO-SQL EVALUATION 001 002 WITH FORMAL VERIFICATION 003 004

005 **Anonymous authors**

006 Paper under double-blind review

007 008 ABSTRACT 009

010 Community-driven Text-to-SQL evaluation platforms play a pivotal role in tracking
011 the state of the art of Text-to-SQL performance. The reliability of the evaluation
012 process is critical for driving progress in the field. Current evaluation methods are
013 largely test-based, which involves comparing the execution results of a generated
014 SQL query and a human-labeled ground-truth on a static test database. Such an
015 evaluation is optimistic, as two queries can coincidentally produce the same output
016 on the test database while actually being different. In this work, we propose a new
017 alternative evaluation pipeline, called SPOTIT, where a formal bounded equivalence
018 verification engine actively searches for a database that differentiates the generated
019 and ground-truth SQL queries. We develop techniques to extend existing verifiers
020 to support a richer SQL subset relevant to Text-to-SQL. A performance evaluation
021 of ten Text-to-SQL methods on the high-profile BIRD dataset suggests that test-
022 based methods can often overlook differences between the generated query and the
023 ground-truth. Further analysis of the verification results reveals a more complex
024 picture of the current Text-to-SQL evaluation.

025 026 1 INTRODUCTION 027

028 Text-to-SQL is one of the fundamental building blocks for designing natural language (NL) interfaces
029 that enable users to access and analyze structured data sources. Translating human questions into
030 executable database queries bridges the gap between non-technical users and complex data systems.
031 This functionality underpins modern chatbots and smart assistants across a wide range of industrial
032 applications, such as observability platforms for monitoring system health (BitsAI, 2025; Splunk,
033 2025), critical business processes (Amazon, 2025), and healthcare (Amazon Web Services, 2024).

034 Due to its practical relevance for commercial products, Text-to-SQL has recently attracted significant
035 attention, leading to the development of a wide range of solutions (Shi et al., 2024). New Text-to-SQL
036 frameworks are announced regularly, and thanks to community-driven evaluation platforms such
037 as BIRD (Li et al., 2024) and Spider (Lei et al., 2024), their performance can be benchmarked and
038 compared in near real time. Given the pivotal role these platforms play in tracking the state of the art,
039 the reliability of their evaluation processes is crucial for driving progress in the field.

040 In this paper, we take a close look at the evaluation process for the accuracy of Text-to-SQL methods.
041 Currently, the process usually involves checking whether the SQL queries generated by a method
042 produce results equivalent to those of the *gold SQLs* (i.e., human-written ground-truth SQLs), under
043 a pre-defined notion of equivalence. Most state-of-the-art evaluation frameworks (Li et al., 2024; Lei
044 et al., 2024) perform this equivalence check through *testing*: executing both queries on a static test
045 database and comparing the results. If the results match, the generated SQL is labeled as correct.
046 Although widely used in practice, the testing-based approach has clear limitations. Because the
047 check is performed on a single database, two different SQL queries may appear equivalent by chance,
048 purely due to the specific data contained in that database. This raises an important question: when
049 the test-based approach marks a generated SQL as correct, how often does it truly produce the same
050 results as the gold SQL in general? The next broader question is: to what extent can the current
051 evaluation process accurately measure the performance of Text-to-SQL methods?

052 We investigate these questions by exploring an alternative correctness evaluation methodology.
053 Instead of relying on test databases to assess equivalence, we propose to actively *search* for databases

that can differentiate the generated SQL from the gold SQL. The search-based evaluation naturally provides stronger correctness guarantees and enables a more rigorous measurement of accuracy. Since providing complete equivalence guarantee is in general undecidable, we perform SMT-based bounded verification (He et al., 2024), which searches for differentiating databases with specified sizes. We develop a new Text-to-SQL evaluation workflow, SPOTIT, on top of those verification techniques. We significantly extend these techniques to support a new set of SQL operators over strings and dates which are commonly used for Text-to-SQL benchmarks.

Experiments on ten state-of-the-art Text-to-SQL methods on the popular BIRD dataset (Li et al., 2024) suggest that the reported accuracy of these methods drops by 11.3%–14.2% when switching from the official test-based evaluation to SPOTIT. The varying levels of decrease in absolute precision also lead to substantial changes in the order of ranking of the Text-to-SQL methods. Moreover, SPOTIT produces minimal differentiating databases, which enables us to pinpoint the sources of inconsistencies between the generated and gold SQLs. Analysis of these databases uncovers several shortcomings of the current Text-to-SQL evaluation process. Most surprisingly, we find that when the predicted SQL disagrees with the ground truth, it is often the gold SQL that is incorrect.

To summarize, our contributions include:

- SPOTIT, a new evaluation pipeline for Text-to-SQL powered by formal equivalence verification;
- novel SMT-encoding for a set of SQL operators over strings and dates, and proof of its correctness;
- practical strategies for the efficient deployment of SPOTIT;
- a large-scale evaluation of ten state-of-the-art Text-to-SQL methods on the BIRD dataset, which reveals several potential shortcomings of current Text-to-SQL evaluation.

2 PRELIMINARIES

We provide background on Text-to-SQL and formal equivalence checking. Due to space limitation, an overview of related work is present in App. A.

Text-to-SQL problem statement. Given a natural language query N and a database D with schema S , the goal of Text-to-SQL is to map (N, D) to an SQL query Q , such that executing Q on D , denoted $Q(D)$, produces an output relation (table) that answers N .

Text-to-SQL evaluation. The main evaluation mechanism for a Text-to-SQL framework relies on a gold SQL query produced by a human annotator. Hence, for each natural language query N over a database, there exists a gold SQL query Q that represents the human-labelled ground truth of translating N into SQL. Given a generated SQL P and the corresponding gold query Q , current evaluation performs the following check:

$$\text{EX-TEST}(P, Q, D_{\text{test}}) = \begin{cases} 1, & \text{if } \forall r. r \in P(D_{\text{test}}) \leftrightarrow r \in Q(D_{\text{test}}) \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where D_{test} is a test database provided by the benchmark set, and r denotes a row in the result table. In words, EX-TEST compares whether the two tables, $P(D_{\text{test}})$ and $Q(D_{\text{test}})$, contain the same set of rows. In order to more rigorously analyze the equivalence between P and Q , we use formal verification to search for a differentiating database D_{cex} such that $\text{EX-TEST}(P, Q, D_{\text{cex}}) = 0$.

Bounded SQL equivalence checking. Given two SQL queries Q_1 and Q_2 over a schema S and an upper bound K on the relation size, the problem of bounded equivalence checking is to decide whether Q_1 and Q_2 are equivalent, denoted $Q_1 \simeq_{S, K} Q_2$, for all databases D conforming to S such that each relation in D has at most K tuples. Formally,

$$Q_1 \simeq_{S, K} Q_2 \stackrel{\text{def}}{=} \forall D \in \text{Instances}(S). \forall R \in \text{Relations}(D). |R| \leq K \Rightarrow Q_1(D) = Q_2(D),$$

where $\text{Instances}(S)$ represents all database instances conforming to S , and $\text{Relations}(D)$ represents all relations in D . In general, the goal is either to prove the bounded equivalence holds, or to find a counterexample database D_{cex} that disproves the equivalence. Compared with unbounded equivalence checking, which is generally undecidable (Mohamed et al., 2024), bounded equivalence checking can handle a more expressive SQL subset and is guaranteed to uncover small counterexamples (if they exist). These features make bounded verification suitable for large-scale Text-to-SQL evaluation.

VERIEQL. VERIEQL (He et al., 2024) is a recently proposed bounded equivalence checker for SQL queries and, to the best of our knowledge, supports the most expressive subset of SQL among

```

108 N1: "Which is the youngest patient with an abnormal
109 anti-ribonuclear protein level?
110 Please list his or her date of birth."
111 /*Gold SQL Q*/:
112 SELECT T1.birthday
113 FROM patient AS T1
114 INNER JOIN laboratory AS T2
115 ON T1.ID = T2.ID
116 WHERE T2.rnp != '-' OR '+-'
117 ORDER BY T1.birthday DESC LIMIT 1
118 /*Generated SQL P*/:
119 SELECT patient.birthday
120 FROM patient
121 INNER JOIN laboratory
122 ON patient.ID = laboratory.ID
123 WHERE NOT laboratory.rnp IN ('-', '+-')
124 ORDER BY patient.birthday
125 DESC LIMIT 1

```

N₂: "How many male patients who underwent testing between 1995 and 1997 and were subsequently diagnosed with Behcet disease did not stay in the hospital for treatment?"

```

/*Gold SQL Q*/:
SELECT COUNT(T1.id) FROM patient AS T1
INNER JOIN examination AS T2 ON T1.id = T2.id
WHERE T2.diagnosis = 'Behcet' AND T1.sex = 'M'
AND STRFTIME('%Y', T2.examination_date)
BETWEEN '1995' AND '1997' AND T1.admission = '-';
/*Generated SQL P*/:
SELECT COUNT(DISTINCT patient.id)
FROM patient INNER JOIN examination
ON patient.id = examination.id
WHERE patient.sex = 'M' AND
examination.examination_date
BETWEEN '1995-01-01' AND '1997-12-31',
AND examination.diagnosis = 'Behcet'
AND patient.admission = '-';

```

Figure 1: Examples of cases where the generated SQL produces the same output as the gold SQL on the BIRD’s official test database, but SPOTIT finds a database that differentiates the the queries. The parts that explain the mismatch are highlighted. For N_1 , the gold SQL is incorrect. And for N_2 , both SQL queries can be right depending on the interpretation of the NL question.

existing tools. It reduces the verification task to a satisfiability problem by encoding the symbolic execution of the two SQL queries and the *non-equivalence* of the execution results as a satisfiability modulo theories (SMT) formula (Barrett & Tinelli, 2018), which can be solved by an off-the-shelf SMT solver (De Moura & Bjørner, 2008). The bounded equivalence property holds if and only if the formula is unsatisfiable, which means it is not possible to find a database that result in different execution results. Otherwise, a satisfying interpretation of the formula can be decoded to a counterexample database. We significantly extend VERIEQL to support our verification use cases.

3 MOTIVATING EXAMPLES

Before we describe our new verification-based evaluation pipeline, we first discuss main sources of mismatches between the gold SQL and the generated SQL in Text-to-SQL evaluation. There are three main such sources: (1) NL query N is ambiguous, so both the gold and generated SQL queries are justifiable interpretations; (2) N is unambiguous, but the gold SQL query is incorrect (gold SQLs are created manually and thus prone to human errors); (3) N is unambiguous, the gold SQL query is correct, but the generated SQL query is incorrect. Our framework focuses on checking equivalence between the gold SQL and the generated SQL, treating the latter as the best-effort, semantically correct formalization of N . We show that SPOTIT can successfully detect incorrect generated SQLs that are overlooked by existing test-based evaluation. Perhaps more surprisingly, SPOTIT also allows us to spot the first and second sources of mismatch. Fig. 1 shows two illustrative examples.

Example 3.1. Consider the query N_1 : "Which is the youngest patient with an abnormal anti-ribonuclear protein level? Please list his or her date of birth." together with the gold and generated SQL queries. On the development database that BIRD provides, both queries return "1989-08-28". However, SPOTIT found a database on which these two queries are not equivalent (Appendix D.1). In fact, we observe that all ten frameworks that we tested generated SQLs that are not equivalent to the gold query. Upon closer inspection, we find that the gold query is incorrect: its WHERE clause is equivalent to $T2.rnp != '-' \text{ OR } FALSE$, as a string literal like $'+-'$ is interpreted as FALSE in a boolean context, which is not the intended behavior. \square

Example 3.2. Consider another query N_2 : "How many male patients who underwent testing between 1995 and 1997 and were subsequently diagnosed with Behcet disease did not stay in the hospital for treatment?" together with the gold and generated SQL queries. These two queries both return "2" on the BIRD test database. However, the two queries are clearly not equivalent (id is not a primary key of the examination table therefore duplicates are allowed): the generated query counts all examinations per patient, whereas the gold query counts only distinct patients. SPOTIT easily found a database that differentiate the two queries (Appendix D.2). Note that depending on the interpretation of the question, both SQL queries can be correct: the gold SQL can be reasonable if the goal is to understand the hospital workload, while the generated SQL can be reasonable if the goal is to understand the number of unique patients. Hence, we conclude that N_2 is ambiguous. \square

```

162
163     Query  $Q_r$  ::=  $Q$  |  $\text{OrderBy}(Q, \vec{E}, b)$ 
164     Subquery  $Q$  ::=  $R$  |  $\Pi_L(Q)$  |  $\sigma_\phi(Q)$  |  $\rho_R(Q)$  |  $Q \oplus Q$  |  $\text{Distinct}(Q)$  |  $Q \otimes Q$  |  $\text{GroupBy}(Q, \vec{E}, L, \phi)$  |  $\text{With}(\vec{Q}, \vec{R}, Q)$ 
165     Attr List  $L$  ::=  $id(A)$  |  $\rho_a(A)$  |  $L, L$ 
166     Attr  $A$  ::=  $\text{Cast}(\phi)$  |  $E$  |  $\mathcal{G}(E)$  |  $A \diamond A$ 
167     Pred  $\phi$  ::=  $b$  |  $\text{Null}$  |  $A \odot A$  |  $\text{IsNull}(E)$  |  $\vec{E} \in \vec{v}$  |  $\vec{E} \in Q$  |  $\phi \wedge \phi$  |  $\phi \vee \phi$  |  $\neg \phi$ 
168           | PrefixOf( $s, E$ ) | SuffixOf( $s, E$ ) | Like( $s, E$ ) | Contain( $s, E$ )
169     Expr  $E$  ::=  $a$  |  $v$  |  $E \diamond E$  |  $\text{ITE}(\phi, E, E)$  |  $\text{Case}(\vec{\phi}, \vec{E}, E)$  | SubStr( $E_1, E_2, E_3$ ) | Concat( $E_1, E_2$ )
170           | Strftime( $\kappa, E$ ) | JulianDay( $E$ ) | DateShift( $E, i, \delta$ ) | ToInt( $E$ ) | ToDate( $E$ ) | ToStr( $E$ )
171     Join Op  $\otimes$  ::=  $\times$  |  $\bowtie_\phi$  |  $\bowtie_{\phi}$  |  $\bowtie_\phi$  |  $\bowtie_{\phi}$ 
172     Collection Op  $\oplus$  ::=  $\cup$  |  $\cap$  |  $\setminus$  |  $\uplus$  |  $\oplus$  |  $\ominus$ 
173     Arith Op  $\diamond$  ::=  $+$  |  $-$  |  $\times$  |  $/$  |  $\%$ 
174     Logic Op  $\odot$  ::=  $\leq$  |  $<$  |  $=$  |  $\neq$  |  $>$  |  $\geq$ 
175
176      $R \in \text{Relation Names}$   $a \in \text{Attribute Names}$   $v \in \{\text{Null}\} \cup \text{Integers} \cup \text{Dates} \cup \text{Strings}$   $b \in \text{Bools}$   $i \in \text{Integers}$ 
177      $s \in \text{Strings}$   $\mathcal{G} \in \{\text{Count}, \text{Min}, \text{Max}, \text{Sum}, \text{Avg}\}$   $\kappa \in \{\%\text{Y}, \%\text{M}, \%\text{d}\}$   $\delta \in \{\text{Year}, \text{Month}, \text{Day}\}$ 

```

Figure 2: Extended syntax of SQL Queries. New features are in bold.

Note that these examples were overlooked by existing test-based evaluations. On the other hand, using SPOTIT, we found that undetected cases like those are quite common in the BIRD dataset.

4 METHODOLOGY

In this section, we introduce new SMT-encodings for a number of SQL operators over string and date types that were not supported by existing bounded equivalence verification methods but frequently appear in Text-to-SQL benchmarks. Then we present our verification-based evaluation pipeline SPOTIT and discuss practical implementation strategies.

4.1 EQUIVALENCE CHECKING FOR SQL QUERIES

To understand our extension, let us first walk through Example 4.1 to understand how equivalence checking can be encoded as an SMT formula in a verifier like VERIQL (He et al., 2024).

Example 4.1. Consider a schema $\mathcal{S} = \{R \mapsto \{id: \text{int}, dob: \text{date}\}\}$ and the following two queries:

```
Q1=SELECT  $id$  FROM  $R$  WHERE  $id > 1$       Q2=SELECT  $id$  FROM  $R$  WHERE  $id > 2$ 
```

We describe how to encode equivalence checking for a bound (K) of 1 as an SMT formula. First, variables are introduced to represent the database and the execution results. This includes a symbolic database $D = \{R \mapsto [t_1]\}$, where $t_1 = [x_1, x_2]$ is a tuple in R , and x_1, x_2 are integer variables. In addition, tuples $t_2 = [x_3]$ and $t_3 = [x_4]$, are introduced to encode query results: $Q_1(D) = [t_2]$ and $Q_2(D) = [t_3]$, where x_3, x_4 are both integer variables. Note that the number of tuples in R is equal to the bound K . Also note that a date (x_2) is represented as an integer, which is sufficient here but not in general. We later introduce precise encoding of date to support richer operations.

We now describe the constraints over the variables. The first set of constraints ensures that t_2 and t_3 correctly capture the semantics of Q_1 and Q_2 . In this case, t_2 tuple is constrained by $\Phi_{Q_1} = (x_1 > 1 \rightarrow (x_3 = x_1 \wedge \neg \text{Del}(t_2))) \wedge (x_1 \leq 1 \rightarrow \text{Del}(t_2))$, where Del is an uninterpreted function denoting the non-existence of a symbolic tuple. The formula Φ_{Q_1} ensures that only interpretations satisfying $x > 1$ can populate a concrete tuple; otherwise, Q_1 's result is empty. Similarly, t_3 is constrained by $\Phi_{Q_2} = (x_1 > 2 \rightarrow (x_4 = x_1 \wedge \neg \text{Del}(t_3))) \wedge (x_1 \leq 2 \rightarrow \text{Del}(t_3))$.

The second set of constraints encodes that $Q_1(D)$ and $Q_2(D)$ returns different results. In this case, it is simply $t_2 \neq t_3$. The full encoding is a conjunction of all constraints: $\Phi_{Q_1} \wedge \Phi_{Q_2} \wedge (t_2 \neq t_3)$, whose satisfiability can be checked by an SMT solver. A satisfying interpretation to this conjunction corresponds to a database instance that differentiates Q_1 and Q_2 . For example, the queries are not equivalent under the interpretation $\mathcal{I} = \{x_1 \mapsto 2\}$. \square

Extension in SQL encoding. Existing bounded SQL equivalence checker still lacks support for several important features, including precise encoding of dates and strings, which are highly relevant in Text-to-SQL applications. Furthermore, SQL supports computations across many different data types with implicit type casting (e.g., $1 + \text{a}$ and $\text{date}(\text{“2000-01-01”}) + \text{“1”}$), which poses significant challenges to establish precise semantics and encodings. To address these limitations and challenges, we introduce techniques to support dates and strings, along with their manipulations, in the SQL

equivalence checker VERIEQL. We also introduce type conversions across Null, integers, dates, and strings for implicit type casting. For example, in the gold SQL for N_2 (Fig. 1), the output of the STRFTIME function is implicitly converted from a date to an integer.

Fig. 2 presents our supported SQL grammar. Specifically, the query language introduces type conversions among various data types (e.g., $\text{ToInt}(E)$, $\text{ToDate}(E)$, and $\text{ToString}(E)$), which allows us to precisely establish the semantics of dates and strings and enhances the expressiveness of our SQL subset. We also incorporate additional expressions and predicates for data and string manipulations, such as date formatting $\text{Strftime}(\kappa, E)$, Julian day $\text{JulianDay}(E)$, string pattern matching $\text{PrefixOf}(s, E)$, $\text{SuffixOf}(s, E)$, $\text{Like}(s, E)$, and string truncation $\text{SubStr}(E_1, E_2, E_3)$. The symbolic encoding for these extended expressions and predicates is formally presented in Appendix F.

As an example, we describe how to precisely encode a date variable, which is very common in Text-to-SQL. For instance, the date of birth and the time of a transaction are naturally modeled with the date type. Previously, date was encoded as a single integer variable (see Example 4.1). Although this coarse representation still enables the encoding of certain date operations (e.g., comparison), it does not necessarily support all date operations, such as date-formatting, which is used in the gold SQL query for N_2 in Fig. 1. As a date can be viewed as a triplet (year, month, day), we introduce three integer variables y , m , and d , and constrain their values with the following formula Φ :

$$\begin{aligned} \Phi = \Phi_1 \wedge \Phi_2 \wedge \Phi_3, \text{ where } \Phi_1 = \text{MIN_YEAR} \leq y \leq \text{MAX_YEAR}, \Phi_2 = 1 \leq m \leq 12, \\ \Phi_3 = 1 \leq d \wedge (\vee_{c \in \{1,3,5,7,8,10,12\}} m = c \rightarrow d \leq 31) \\ \wedge (m = 2 \rightarrow d \leq 28 + \text{ite}(\text{leap}(y), 1, 0)) \wedge (\vee_{c \in \{4,6,9,11\}} m = c \rightarrow d \leq 30) \end{aligned}$$

The term $\text{leap}(y)$ encodes the leap year condition: $y \% 4 = 0 \wedge (y \% 100 \neq 0 \vee y \% 400 = 0)$. Constraints Φ_1 , Φ_2 , and Φ_3 restrict the possible values of the year, the month, and the day, respectively. For example Φ_1 specifies the valid range of the year, which is specific to the database engine. For example, SQLite only accepts dates between “0000-01-01” and “9999-12-31”; in which case MIN_YEAR is 0 and MAX_YEAR is 9999. This refined representation allows us to precisely encode a rich set of date operations and analyze more SQL queries compared to the previous encoding.

Equivalence under set semantics. SQL equivalence checkers typically support equivalence under bag semantics and list semantics. However, some Text-to-SQL evaluation platforms, such as BIRD (Li et al., 2024), by default adopt equivalence under set semantics (see equation 1). This can be expressed as an SMT constraint. Given two query results with symbolic tables $R_1 = [t_1, \dots, t_n]$ and $R_2 = [r_1, \dots, r_m]$, the condition that R_1 and R_2 are equivalent under set semantics is as follows:

$$\bigwedge_{i=1}^n (\neg \text{Del}(t_i) \rightarrow \bigvee_{j=1}^m (\neg \text{Del}(r_j) \wedge t_i = r_j)) \wedge \bigwedge_{j=1}^m (\neg \text{Del}(r_j) \rightarrow \bigvee_{i=1}^n (\neg \text{Del}(t_i) \wedge r_j = t_i)) \quad (2)$$

On a high level, equivalence is defined by mutual set containment: $R_1 = R_2$ iff $R_1 \subseteq R_2$ and $R_2 \subseteq R_1$. But since some tuples might be deleted due to WHERE clauses, we restrict set containment to non-deleted tuples, i.e., those satisfying $\neg \text{Del}(t)$.

Correctness of the encodings. We now state the correctness of our symbolic encoding for the extended expressions and predicates, as well as the equivalence under set semantics. Proof of these theorems is in Appendix G. As we encode the symbolic execution of queries, to prove the correctness of our approach, we need to show that our symbolic execution coincides with the concrete execution. This involves showing that given an expression E , the satisfying interpretation of E ’s symbolic execution result is identical to the concrete execution result of E . Thm. 1 states that formally.

Theorem 1 (Correctness of expression encoding). *Let D be a database over schema \mathcal{S} , xs be a tuple list, and E be an expression. Consider a symbolic database Γ over \mathcal{S} , a list of symbolic tuples \mathcal{T} , and E ’s symbolic encoding $\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$. For any satisfying interpretation \mathcal{I} with $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs$, evaluating the expression E over the database D and the tuple list xs yields the interpretation of E ’s symbolic encoding $\mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$, i.e., $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs \Rightarrow \llbracket E \rrbracket_{D, xs} = \mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$.*

Similarly, given a predicate ϕ , the satisfying interpretation of ϕ ’s symbolic execution result is also identical to the concrete execution result of ϕ . This is formally stated in Appendix G. Lastly, we state the correctness of our encoding for equivalence under set semantics.

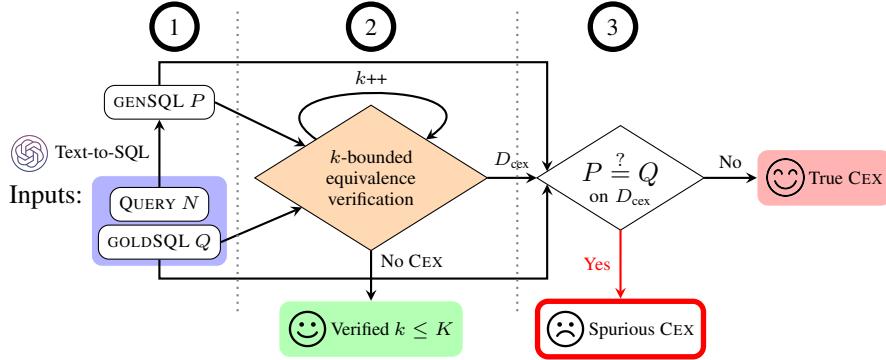
Theorem 2 (Equivalence under set semantics). *Given two relations $R_1 = [t_1, \dots, t_n]$ and $R_2 = [r_1, \dots, r_m]$, if formula (2) is valid, then R_1 and R_2 are equivalent under set semantics.*

270 **Algorithm 1** Bounded equivalence checking
 271 **Require:** Database schema \mathcal{S} , gold SQL query Q , generated SQL
 272 query P , time limit T , bound K
 273 **Ensure:** A counterexample D_{cex}
 274 1: **function** EQUIVCHECK(\mathcal{S}, Q, P, T, K)
 275 2: **for** $k \in [1, K]$ **do**
 276 3: $\text{res}, D_{\text{cex}} \leftarrow \text{CHECKBOUND}(\mathcal{S}, P, Q, k, T)$
 277 4: **if** $\text{res} = \text{EQUIVALENT}$ **then continue**
 278 5: **▷** Bounded equivalence under k
 279 6: **else if** $\text{res} = \text{NON-EQUIVALENT}$ **then**
 280 7: **▷** Find a counterexample
 281 8: **▷** Validate the counterexample on the backend DBMS
 282 9: **if** $\neg \text{EX-TEST}(P, Q, D_{\text{cex}})$ **then**
 283 10: **return** $\{D_{\text{cex}}\}$
 284 11: **else break** **▷** Timeout, unsupported, undecidable queries
 285 12: **return** \emptyset

286 **Algorithm 2** SPOTIT⁺
 287 **Require:** Database \mathcal{S} , user query N , gold SQL query Q , Text-to-
 288 SQL frameworks \mathcal{M} , time limit T and bound K
 289 **Ensure:** Counterexamples D_{cexs}
 290 1: **function** SPOTIT⁺($\mathcal{S}, N, \mathcal{M}, T, K$)
 291 2: $D_{\text{cexs}} \leftarrow \emptyset$
 292 3: **for** $m \in \mathcal{M}$ **do**
 293 4: $P \leftarrow m(\mathcal{S}, N)$ **▷** Generate SQL query P using m
 294 5: $D_{\text{cexs}}[m] \leftarrow \text{EQUIVCHECK}(\mathcal{S}, Q, P, T, K)$
 295 6: **▷** Performing cross-referencing counterexamples
 296 7: $D_{\text{cexs}}^* \leftarrow \bigcup_{m \in \mathcal{M}} D_{\text{cexs}}[m]$
 297 8: **for** $m \in \mathcal{M}$ **do**
 298 9: **for** $D \in D_{\text{cexs}}^* \setminus D_{\text{cexs}}[m]$ **do**
 299 10: **if** $\neg \text{EX-TEST}(P, Q, D)$ **then**
 300 11: $D_{\text{cexs}}[m] \leftarrow D_{\text{cexs}}[m] \cup \{D\}$
 301 12: **return** D_{cexs}

4.2 SPOTIT: A SEARCH-BASED TEXT-TO-SQL EVALUATION PIPELINE

283 Fig. 3 presents a high-level workflow of our approach that consists of three conceptual phases.



299 Figure 3: Three main phases of SPOTIT.

300 **① Input phase.** Given a NL question N and its corresponding gold SQL query Q , a Text-to-SQL
 301 framework takes as input N and generates a SQL query P . Both Q and P are passed to phase ②.

303 **② Verification phase.** The goal is to find a counterexample database instance on which the queries
 304 Q and P produce different outputs. For a given bound $k \leq K$, we perform bounded equivalence
 305 checking between Q and P . If the queries are proved equivalent, then we increase k by one for
 306 the next verification check. Furthermore, we cannot find any counterexample under all bounds and
 307 conclude that they are verified up to the bound k . On the other hand, if the queries are proved to be
 308 non-equivalent under some bound, we proceed to phase ③ for a further validation of D_{cex} .

309 **③ Validation phase.** Given the queries Q and P and a counterexample D_{cex} returned by verification
 310 algorithm, we must verify that this counterexample is non-spurious. There are two main reasons
 311 spurious counterexamples can arise in the verification engine. Either because some operators are
 312 over-approximated in the SMT encoding or the SQL query admits non-deterministic behaviors that
 313 cannot be modeled. Therefore, we execute the queries on the counterexample database (e.g., in
 314 SQLite) and check whether the results actually differ. D_{cex} is viewed valid if the results remain
 315 different; otherwise, we report this spurious case to the developers.

316 Alg.1 implements the second and third phases. For a given bound $k \leq K$, it first checks bounded
 317 equivalence between Q and P (line 3). If the queries are proven to be non-equivalent (line 6) under
 318 some bound, we validate that the counterexample database is indeed a true counterexample (line 9)
 319 and return it if this is the case. If the queries are proven to be equivalent in line 3, then we increase k
 320 by one for the next verification step. If the verifier cannot find any counterexample under all bounds,
 321 Alg.1 returns an empty set. Finally, if the verifier times out on a bound k , or the query is unsupported
 322 or undecidable, it also returns an empty set.

323 **Cross-checking counter-examples.** One observation we make is that as we progress through the
 frameworks, we collect a set of counterexamples that separate the gold query from the generated

queries. Hence, we realized that these counterexamples can be reused as checks across all frameworks, as they might generalize across frameworks. Alg.2 implements this idea. First, it obtains counterexample databases, if they exist, for all frameworks by calling Alg.1 (lines 3–5). Then, it iterates over all frameworks again and tests equivalence between Q and P on these counterexample databases (lines 7–11). Empirically, this improves the effectiveness of our approach.

5 EXPERIMENTAL EVALUATION

In this section, we investigate the effect of using SPOTIT as the evaluation methodology for Text-to-SQL tasks. We are interested in the following questions:

- How much more SQL queries does our extension of VERIEQL support?
- Can SPOTIT provide more rigorous accuracy evaluation than test-based approaches?
- Can SPOTIT reveal shortcomings in existing Text-to-SQL evaluations?

Experimental Setup. We consider all 1,533 question-SQL pairs from the development set of BIRD (Li et al., 2023b), a state-of-the-art dataset for evaluating Text-to-SQL methods. The questions span 11 different databases from different professional domains, such as education, healthcare, and sports. The official BIRD leaderboard¹ contains over 80 Text-to-SQL methods and are updated frequently. Not all methods are open-source or have predictions publicly available. Therefore, we reached out to the developers of top-performing Text-to-SQL frameworks on the BIRD leaderboard and obtained the generated SQL queries for 10 of them, which constitutes a representative subset of state-of-the-art Text-to-SQL methods. The methods are listed in Tab. 1.

We first evaluate the predictions of each method using BIRD’s official test-based execution accuracy metric (EX-TEST), which, as described in Eq. 1, compares the results of executing the generated and gold queries on a given test database. For predictions that are deemed correct by EX-TEST, we apply SPOTIT to perform a more rigorous analysis. We implemented SPOTIT on top of VERIEQL (He et al., 2024), which we extended using the methods described in Sec. 4.1. To generate practically relevant counterexamples, we also extend the verification condition to exclude degenerate counterexamples that result in empty for one SQL and NULL for the other SQL.

We consider three variants of SPOTIT: (i) SPOTIT: Alg. 2 instantiated with the extended verification engine but without cross-checking (lines 7–11); (ii) SPOTIT⁻: Alg. 2 instantiated with vanilla VERIEQL and without cross-checking; (iii) SPOTIT⁺: Alg. 2 with cross-checking. We verify each generated-gold SQL pair up to a bound (K) of 5. Each verifier call is given one physical core, 8GB memory, and a CPU timeout of 600 seconds. In practice, a counterexample can typically be found within seconds, as reported below. Experiments were performed on a cluster equipped with Dell PowerEdge R6525 CPU servers featuring 2.6-GHz AMD CPU cores.

Performance of Verification Engine. We first evaluate the effect of our extensions to the original VERIEQL engine (He et al., 2024) in terms of *coverage*, defined as the fraction of generated-gold-SQL pairs that can be encoded into an SMT query. In addition, we measure the average runtime of SPOTIT on questions where a valid differentiating database is found. The results are shown in Tab. 2. Our extensions significantly increase the coverage of the verification engine on relevant questions (i.e., ones deemed correct by EX-TEST) for each method, allowing us to formally analyze a larger number of generated SQL queries. For example, for CSC-32B, the coverage increases from 84.83% to 94.88%, which corresponds to 110 additional supported questions $((94.88\% - 84.83\%) * 1094)$.

The average time taken by SPOTIT to find a counterexample is under 4 seconds for all methods, which, combined with the fact that the analysis for each question can be done in parallel, confirms that SPOTIT is already a practical method for formally comparing generated SQLs with gold SQLs.

Table 1: The Text-to-SQL methods we evaluated

| Entry | Acronym |
|--|----------|
| Alpha-SQL + Qwen2.5-Coder-32B (Li et al., 2025a) | ALPHA |
| CSC-SQL + Qwen2.5-Coder-7B (Sheng & Xu, 2025a) | CSC-7B |
| CSC-SQL + XiYanSQL (Sheng & Xu, 2025a) | CSC-32B |
| GenaSQL-1 (Dönder et al., 2025) | GENA-1 |
| GenaSQL-2 (Dönder et al., 2025) | GENA-2 |
| RSL-SQL + GPT-4o (Cao et al., 2024) | RSL |
| OmniSQL-32B(Li et al., 2025b) | OMNI-MAJ |
| GSR (anonymous authors) | GSR |
| CHESS _{IR+CG+UT} (Talaei et al., 2024a) | CHESS |
| SLM-SQL + Qwen2.5-Coder-1.5B Sheng & Xu (2025c) | SLM |

¹<https://bird-bench.github.io/>

378
 379
 380
 381
 382
Table 2: % of SQL pairs supported by SPOTIT⁻ and SPOTIT. For SPOTIT, also the average time in seconds on pairs where CEXs are found by the verifier and how many of them are non-spurious.

| Method (# quest.) | SPOTIT ⁻ (%) | SPOTIT (%) | Avg. Time | Valid. (%) |
|-------------------|-------------------------|------------|-----------|--------------|
| ALPHA (1064) | 84.87 | 93.89 | 3.10 | 96.15 |
| CHESS (976) | 87.40 | 97.13 | 1.40 | 93.34 |
| CSC-32B (1094) | 84.83 | 94.88 | 3.24 | 94.46 |
| CSC-7B (1061) | 85.77 | 96.14 | 3.93 | 95.10 |
| GENA-1 (1062) | 84.56 | 94.92 | 1.01 | 94.52 |
| GENA-2 (1082) | 84.47 | 94.55 | 0.93 | 95.42 |
| GSR (1020) | 84.51 | 93.63 | 1.12 | 94.86 |
| OMNI-MAJ (1026) | 86.65 | 95.61 | 1.36 | 95.83 |
| RSL (1038) | 86.03 | 95.18 | 1.64 | 95.62 |
| SLM (973) | 85.92 | 94.24 | 1.36 | 95.05 |

390
 391
 392
 393
 394
Table 3: Performance of Text-to-SQL methods using EX-TEST, EX-SPOTIT, and EX-SPOTIT⁺ on the 1533 BIRD-dev benchmarks.

| | EX-TEST | | EX-SPOTIT | | EX-SPOTIT ⁺ | |
|----------|----------|------|-----------|------|------------------------|------|
| | Acc. (%) | Rank | Acc. (%) | Rank | Acc. (%) | Rank |
| CSC-32B | 71.32 | 1 | 58.80 | 3 | 57.82 | 4 |
| GENA-2 | 70.53 | 2 | 59.84 | 1 | 59.13 | 1 |
| ALPHA | 69.36 | 3 | 55.87 | 6 | 55.02 | 6 |
| GENA-1 | 69.23 | 4 | 59.45 | 2 | 59.00 | 2 |
| CSC-7B | 69.17 | 5 | 58.54 | 4 | 57.95 | 3 |
| RSL | 67.67 | 6 | 56.58 | 5 | 55.80 | 5 |
| OMNI-MAJ | 66.88 | 7 | 54.69 | 7 | 54.04 | 7 |
| GSR | 66.49 | 8 | 54.56 | 8 | 53.72 | 8 |
| CHESS | 63.62 | 9 | 52.87 | 9 | 52.35 | 9 |
| SLM | 63.43 | 10 | 51.37 | 10 | 50.98 | 10 |

395
 396
 397
 398
 399
 400
 401
 402
 403
 404
 405
 406
 407
 Moreover, a high percentage (up to 96.15%) of the counterexamples found by the verifier are successfully validated, which suggests that our SMT encoding is sufficiently precise in practice.

408
 409
 410
 411
 412
 413
 414
 415
Comparing test-based evaluation with SPOTIT. We now evaluate the accuracy of each Text-to-SQL method based on EX-TEST, EX-SPOTIT, and EX-SPOTIT⁺. As shown in Tab. 3, the accuracy of each method drops significantly when SPOTIT is used to check query equivalence. For example, the accuracy of CSC-32B drops from 71.32% to 58.80% with SPOTIT, and further to 57.82% when cross-checking is enabled. This means that there are 207 generated SQLs ($1533 * (71.32\% - 57.82\%)$) that passed the test on the official test databases, but were differentiated from the gold SQL by SPOTIT. Overall, SPOTIT resulted in a decrease in accuracy ranging from 9.8% to 13.5%, and cross-checking results in a small further decrease, by up to 1%. Interestingly, the ranking of the Text-to-SQL methods also changes substantially when evaluated under the more stringent, verification-based metrics, particularly in the top half of the table. For example, CSC-32B, which is ranked 1st by the official test-based metric, drops to 4th place when evaluated by SPOTIT⁺. And the 3rd place method ALPHA drops to the 6th place. These results indicate that test-based methods can in many cases overlook differences between the generated SQL and the gold SQL, which might lead to misrepresentation of the actual performance (both *absolute* and *relative*) of existing Text-to-SQL methods.

416
 417
 418
 419
 420
 421
 422
 423
 424
 425
 426
 427
 428
 429
 430
 431
The effect of K. To study the effect of the choice of the verification bound K , we vary its value from 1 to 7 and run SPOTIT on the predictions of CSC-32B, the best model according to EX-TEST. As shown in Fig. 4, SPOTIT was able to find significantly more differentiating databases when K increases from 1 to 2 and the gain is marginal pass $K = 3$. This justifies our choice of $K = 5$.

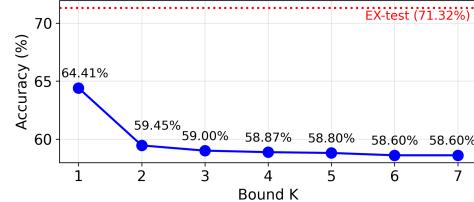


Figure 4: The effect of bound K .

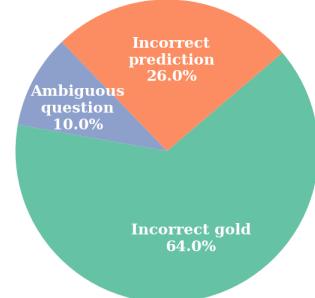


Figure 5: A breakdown of the primary reason for the difference between generated and gold SQLs.

432
 433
 434
 435
 436
 437
 438
 439
 440
 441
Manual inspection of SPOTIT counterexamples. As SPOTIT performs bounded verification, the differentiating databases it finds are guaranteed to be minimal, which makes it easy to analyze them and understand the source of difference between the generated and gold SQLs. We manually examined the counterexamples for a random sample of 50 queries generated by CSC-32B and found that the difference between a generated SQL and a gold SQL can be primarily attributed to the three reasons that we described in Section 3: ambiguous question, incorrect gold SQL, and incorrect generated SQL. Fig. 5 shows a breakdown of the primary attributed reasons for those sampled questions. Surprisingly, while incorrect predictions do constitute a significant portion (26%), more often than not, the gold SQL itself is problematic. There are also a small fraction of cases (10%) where the question itself can be interpreted in multiple ways and therefore admits different answers. We discussed two examples of incorrect gold SQLs and ambiguous questions in Fig. 1. Additional examples of each type of issues, along with the databases found by SPOTIT, are provided and discussed in App. D.

432 **Additional analysis on Spider 2.0 benchmarks.** To assess the generalizability of SPOTIT to more
 433 complex Text-to-SQL tasks, we evaluate it on the
 434 recently introduced Spider 2.0 benchmark (Lei
 435 et al., 2024). We consider OMNISQL (Li et al.,
 436 2025b), a state-of-the-art Text-to-SQL method and
 437 GPT-5² on the 135 SQLite questions. These
 438 methods pass EX-TEST on 46 and 57 queries, respectively, which is competitive with the top entries
 439 on the Spider 2.0 leaderboard.³ As shown in Tab. 4, SPOTIT finds differentiating databases for 16
 440 ((34.1% – 22.2%) * 135) and 8 query pairs deemed correct by test-based evaluation respectively
 441 for OMNISQL and GPT-5. SPOTIT’s runtime for finding counterexamples remains low. We believe
 442 this is due to the fact that the schemas in Spider 2.0 are only moderately larger than those in BIRD
 443 (97.6 vs. 78.6 columns) and counter-examples can usually be detected with small values of K . The
 444 main verification challenge when it comes to Spider 2.0 lies in the number of SQL operators required
 445 for the queries but currently unsupported by our verification engine, which resulted in a smaller
 446 percentage of supported query pairs. Upon closer examination, 52.6% of the unsupported Spider
 447 2.0 queries involve the window function (i.e., `OVER` clauses). Overall, these results indicate that
 448 SPOTIT is also useful for uncovering query discrepancies overlooked by test-based methods on the
 449 challenging Spider 2.0 benchmarks.

450 **Summary of findings and implications.** We now summarize the findings of our evaluation of a
 451 state-of-the-art Text-to-SQL evaluation dataset BIRD using SPOTIT and discuss their implications.

452 *Finding 1: Existing test-based correctness metrics that involve executing the generated SQL and the*
 453 *gold SQL on static test databases can overlook significant variations in output data returned by the*
 454 *generated and gold SQLs.* A search-based evaluation metric, such as SPOTIT, can serve as a practical
 455 alternative that provides additional perspectives on the performance of Text-to-SQL methods.

456 *Finding 2: there is a significant number of problematic gold SQLs in existing Text-to-SQL benchmark*
 457 *sets.* As shown by examples in Tab. 1 and App. D, in many cases, the issue can be hard to detect, yet
 458 can cause significantly different behaviors from the intended one. The presence of incorrect gold
 459 SQLs makes it hard to determine the true optimal performance on a benchmark set, as even a perfect
 460 Text-to-SQL method cannot achieve 100% accuracy.

461 Based on our result analysis for CSC-32B, we speculate that when
 462 most Text-to-SQL methods disagree with the gold SQL, the gold
 463 SQL is likely problematic. To validate this, we count the number
 464 of times that a prediction for a question is deemed correct by
 465 EX-TEST but incorrect by SPOTIT⁺ across all 10 Text-to-SQL
 466 methods. As shown in Fig. 6, there are 36 questions on which all
 467 methods generated queries that differ from the gold SQL. Manual
 468 inspection suggests that 31 of those 36 cases have problematic
 469 gold SQLs, 3 have ambiguous questions, and only 2 represent
 470 genuine errors in the generated SQLs.

471 While so far we have focused on incorrect gold SQLs overlooked
 472 by EX-TEST, our investigation begs the question: *when the generated*
 473 *query differs from the gold SQL, how often in general is the gold SQL*
 474 *problematic?* Fig. 7 shows the number of times the
 475 prediction for a question is deemed incorrect by EX-TEST across
 476 the 10 Text-to-SQL methods, for questions where CSC-32B’s
 477 predictions failed EX-TEST. There are 294 questions where at least
 478 8 of the other 9 methods also failed EX-TEST. If shared disagreement
 479 with gold SQL is also a good indicator for problematic gold
 480 SQL in this case, then even a perfect Text-to-SQL method might
 481 not be able to achieve an EX-TEST score much higher than 80%.

Table 4: Evaluation of SPOTIT on Spider 2.0.

| | EX-TEST | | EX-SPOTIT | |
|---------|----------|----------|---------------|-----------|
| | Acc. (%) | Acc. (%) | Supported (%) | Avg. Time |
| OMNISQL | 34.1 | 22.2 | 60.9% | 3.4 |
| GPT-5 | 42.2 | 36.3 | 50.9% | 1.1 |

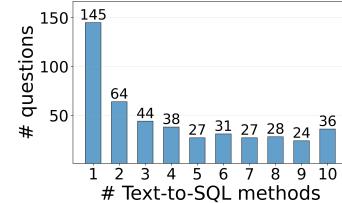
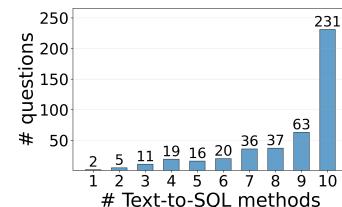
Figure 6: A breakdown of questions that passed EX-TEST but failed SPOTIT⁺.

Figure 7: A breakdown of questions for which CSC-32B’s predictions failed EX-TEST.

²We use the same prompt as used in OMNISQL.³<https://spider2-sql.github.io/>. We were not able to obtain predictions of Text-to-SQL methods on the leaderboard because they are predominantly closed source.

486 on BIRD-dev. As the time of completing this manuscript, the best EX-TEST score for BIRD-dev
 487 achieved by any method on the official leaderboard is 76.14%.
 488

489 Large-scale benchmark sets inevitably contain problematic gold SQLs. Indeed, multiple sources have
 490 found examples of problematic gold SQLs in the BIRD dataset (Hui, 2024; Wretblad et al., 2024),
 491 and some of them have already been addressed by the maintainers. SPOTIT is the first approach that
 492 can provide minimal, easily analyzable databases to differentiate generated and gold SQLs, and can
 493 help to systematically uncover problematic gold SQLs.
 494

495 *Finding 3: A substantial number of questions in the Text-to-SQL dataset can be interpreted in different
 496 ways, thus admitting different SQL queries.* While ambiguity is inherent in natural language, judging
 497 the correctness of a generated SQL query based on a single gold SQL query when the natural language
 498 question admits multiple interpretations might result in unfair penalization of Text-to-SQL methods.
 499

500 *Finding 4 (for the verification community): SMT-based equivalence verification techniques can
 501 already support a large fraction of practical SQL queries.* Our results demonstrate that verification
 502 can often be completed within seconds. Due to the practical relevance of Text-to-SQL, we believe
 503 there is motivation for the verification community to invest more resources to precisely cover a larger
 504 fragment of SQL. In App. C, we discuss further extensions that would be especially useful according
 505 to our evaluation. Another significant next step to incorporate user preferences (potentially from
 506 natural languages) to search for particular types of counterexample databases. One way to achieve
 507 this is to encode the preferences as additional constraints in the SMT formulation of the verification
 508 problem.
 509

6 RELATED WORK

510 Popular evaluation platforms such as BIRD-SQL (Li et al., 2024) and Spider 2.0 (Lei et al., 2024)
 511 evaluate query correctness by testing on predefined database instances. Several additional evaluations
 512 have been proposed to take into account partially correct generated queries (Pinna et al., 2025),
 513 efficiency of query executions (Zhang et al., 2024), and ambiguity in the questions (Li et al., 2023a;
 514 BIRD, 2025). However, the final correctness check is still via testing on a static database. Formal
 515 SQL equivalence checking broadly falls into two categories, full-fledged verification (Chu et al.,
 516 2017c; 2018; Zhou et al., 2022; 2024; Wang et al., 2024) and bounded verification (Veanes et al.,
 517 2010; Chu et al., 2017a;b; He et al., 2024). To the best of our knowledge, VERIEQL (He et al., 2024)
 518 supports the most expressive SQL fragments, while also offering extensibility for new features. We
 519 significantly extend the VERIEQL framework to support date and string types as well as a number of
 520 common operators for the Text-to-SQL evaluation task. Test data generation methods can also be
 521 useful for detecting query non-equivalence (Chandra et al., 2015; Somwase et al., 2024; Zhong et al.,
 522 2020). However, when the counterexamples require very specific structures, random fuzzing/testing
 523 can become unreliable in refuting equivalence. In contrast, SPOTIT systematically searches over the
 524 space of possible differentiating databases, finds minimal counter-examples, and provides a formal
 525 guarantee: if SPOTIT deems two SQL queries equivalent, then no counterexample of size less than a
 526 fixed number exists. A more detailed review of related work can be found in App. A.
 527

7 CONCLUSION

528 We presented SPOTIT, the first verification-based evaluation pipeline for Text-to-SQL. We introduced
 529 techniques to support a richer SQL grammar, which enabled us to efficiently analyze a large fragment
 530 of SQL queries commonly seen in Text-to-SQL tasks. Our initial motivation for developing SPOTIT
 531 was to examine the extent to which the accuracy of a Text-to-SQL method is overestimated by
 532 test-based evaluation, which is widely adopted as the default metric on high-profile Text-to-SQL
 533 evaluation platforms. However, a closer inspection of the verification results revealed a far more
 534 complex picture. While SPOTIT can indeed detect incorrect generated SQL queries that were
 535 overlooked by test-based methods, a significant portion of the inconsistency between the gold and
 536 generated SQLs can be explained by the benchmarks themselves—either due to problematic gold
 537 SQLs or due to ambiguous natural language questions. We discussed the implications of and the
 538 next steps from our findings, and hope that our work will motivate further work on evaluating and
 539 improving Text-to-SQL evaluation frameworks.

540 REFERENCES
541

542 Amazon. Build a text-to-sql solution for data consistency in generative ai using amazon nova,
543 September 2025. URL <https://aws.amazon.com/blogs/machine-learning/build-a-text-to-sql-solution-for-data-consistency-in-generative-ai-using-amazon-nova/>.

544

545 Amazon Web Services. How msd uses amazon bedrock to translate natural language into sql for
546 complex healthcare databases. <https://aws.amazon.com/blogs/machine-learning/how-merck-uses-amazon-bedrock-to-translate-natural-language-into-sql-for-complex-healthcare-databases/>, 2024. Accessed: 2025-09-21.

547

548 Clark Barrett and Cesare Tinelli. Satisfiability modulo theories. In *Handbook of model checking*, pp.
549 305–343. Springer, 2018.

550

551 Team BIRD. Bird-interact: Re-imagine text-to-sql evaluation via lens of dynamic interactions.
552 <https://github.com/bird-bench/BIRD-Interact>, 2025. Accessed: 2025-06-01.

553

554 Team BitsAI. Bits ai:scale your teams with autonomous ai agents across monitoring, development,
555 and security, September 2025. URL <https://www.datadoghq.com/product/platform/bits-ai/>.

556

557 Zhenbiao Cao, Yuanlei Zheng, Zhihao Fan, Xiaojin Zhang, Wei Chen, and Xiang Bai. Rsl-sql: Robust
558 schema linking in text-to-sql generation. *arXiv preprint arXiv:2411.00073*, 2024.

559

560 Bikash Chandra, Bhupesh Chawda, Biplab Kar, KV Maheshwara Reddy, Shetal Shah, and S Sudar-
561 shan. Data generation for testing and grading sql queries. *The VLDB Journal*, 24(6):731–755,
562 2015.

563

564 Shuaichen Chang and Eric Fosler-Lussier. How to prompt llms for text-to-sql: A study in zero-shot,
565 single-domain, and cross-domain settings. *arXiv preprint arXiv:2305.11853*, 2023.

566

567 Shumo Chu, Daniel Li, Chenglong Wang, Alvin Cheung, and Dan Suciu. Demonstration of the
568 colette automated sql prover. In *Proceedings of the 2017 ACM International Conference on
569 Management of Data*, pp. 1591–1594, 2017a.

570

571 Shumo Chu, Chenglong Wang, Konstantin Weitz, and Alvin Cheung. Colette: An automated prover
572 for sql. In *CIDR*, pp. 1–7, 2017b.

573

574 Shumo Chu, Konstantin Weitz, Alvin Cheung, and Dan Suciu. Hottsql: Proving query rewrites with
575 univalent sql semantics. *Acm sigplan notices*, 52(6):510–524, 2017c.

576

577 Shumo Chu, Brendan Murphy, Jared Roesch, Alvin Cheung, and Dan Suciu. Axiomatic foundations
578 and algorithms for deciding semantic equivalences of sql queries. *arXiv preprint arXiv:1802.02229*,
579 2018.

580

581 Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *International conference on
Tools and Algorithms for the Construction and Analysis of Systems*, pp. 337–340. Springer, 2008.

582

583 Yusuf Denizay Dönder, Derek Hommel, Andrea W Wen-Yi, David Mimno, and Unso Eun Seo Jo.
584 Cheaper, better, faster, stronger: Robust text-to-sql without chain-of-thought or fine-tuning. *arXiv
585 preprint arXiv:2505.14174*, 2025.

586

587 Xuemei Dong, Chao Zhang, Yuhang Ge, Yuren Mao, Yunjun Gao, lu Chen, Jinshu Lin, and Dongfang
588 Lou. C3: Zero-shot text-to-sql with chatgpt, 2023.

589

590 Yingqi Gao, Yifu Liu, Xiaoxia Li, Xiaorong Shi, Yin Zhu, Yiming Wang, Shiqi Li, Wei Li, Yuntao
591 Hong, Zhiling Luo, Jinyang Gao, Liyu Mou, and Yu Li. A preview of xiyan-sql: A multi-generator
592 ensemble framework for text-to-sql, 2025. URL <https://arxiv.org/abs/2411.08599>.

593

594 Yang He, Pinhan Zhao, Xinyu Wang, and Yuepeng Wang. Verieql: Bounded equivalence verification
595 for complex sql queries with integrity constraints. *Proc. ACM Program. Lang.*, 8(OOPSLA1),
596 April 2024. doi: 10.1145/3649849. URL <https://doi.org/10.1145/3649849>.

594 Binyuan Hui. Panel of BIRD Annotation Issues. <https://github.com/AlibabaResearch>
 595 /DAMO-ConvAI/issues/39, 2024. Accessed: May 2024.
 596

597 Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin Su, Zhaoqing Suo,
 598 Hongcheng Gao, Wenjing Hu, Pengcheng Yin, et al. Spider 2.0: Evaluating language models on
 599 real-world enterprise text-to-sql workflows. *arXiv preprint arXiv:2411.07763*, 2024.

600 Boyan Li, Jiayi Zhang, Ju Fan, Yanwei Xu, Chong Chen, Nan Tang, and Yuyu Luo. Alpha-sql:
 601 Zero-shot text-to-sql using monte carlo tree search. *arXiv preprint arXiv:2502.17248*, 2025a.
 602

603 Haoyang Li, Shang Wu, Xiaokang Zhang, Xinmei Huang, Jing Zhang, Fuxin Jiang, Shuai Wang,
 604 Tieying Zhang, Jianjun Chen, Rui Shi, et al. Omnisql: Synthesizing high-quality text-to-sql data at
 605 scale. *arXiv preprint arXiv:2503.02240*, 2025b.

606 Jinyang Li, Binyuan Hui, Ge Qu, Binhua Li, Jiaxi Yang, Bowen Li, Bailin Wang, Bowen Qin, Rongyu
 607 Cao, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin Chen-Chuan
 608 Chang, Fei Huang, Reynold Cheng, and Yongbin Li. Can LLM already serve as A database
 609 interface? A big bench for large-scale database grounded text-to-sqls. *CoRR*, abs/2305.03111,
 610 2023a. doi: 10.48550/ARXIV.2305.03111. URL <https://doi.org/10.48550/arXiv.2305.03111>.
 611

612 Jinyang Li, Binyuan Hui, Ge Qu, Jiaxi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin,
 613 Ruiying Geng, Nan Huo, et al. Can ILM already serve as a database interface? A big bench for
 614 large-scale database grounded text-to-sqls. *Advances in Neural Information Processing Systems*,
 615 36:42330–42357, 2023b.

616 Jinyang Li, Binyuan Hui, Ge Qu, Binhua Li, Jiaxi Yang, Bowen Li, Bailin Wang, Bowen Qin, Rongyu
 617 Cao, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin Chen-Chuan
 618 Chang, Fei Huang, Reynold Cheng, and Yongbin Li. BIRD-SQ. <https://bird-bench.git.hub.io/>, 2024. Accessed: Sep 2025.
 619

620 Aiwei Liu, Xuming Hu, Lijie Wen, and Philip S. Yu. A comprehensive evaluation of chatgpt's
 621 zero-shot text-to-sql capability, 2023.

622 Yifu Liu, Yin Zhu, Yingqi Gao, Zhiling Luo, Xiaoxia Li, Xiaorong Shi, Yuntao Hong, Jinyang Gao,
 623 Yu Li, Bolin Ding, and Jingren Zhou. Xiyan-sql: A novel multi-generator framework for text-to-sql,
 624 2025. URL <https://arxiv.org/abs/2507.04701>.
 625

626 Mudathir Mohamed, Andrew Reynolds, Cesare Tinelli, and Clark Barrett. Verifying sql queries using
 627 theories of tables and relations. In *Proceedings of 25th Conference on Logic for Pro*, volume 100,
 628 pp. 445–463, 2024.

629 Giovanni Pinna, Yuriy Perezhoin, Luca Manzoni, Mauro Castelli, and Andrea De Lorenzo. Redefin-
 630 ing text-to-sql metrics by incorporating semantic and structural similarity. *Scientific Reports*, 15,
 631 July 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-04890-9.
 632

633 Juan Sequeda, Dean Allemand, and Bryon Jacob. A benchmark to understand the role of knowledge
 634 graphs on large language model's accuracy for question answering on enterprise sql databases,
 635 2023.

636 Lei Sheng and Shuai-Shuai Xu. Csc-sql: Corrective self-consistency in text-to-sql via reinforcement
 637 learning. *arXiv preprint arXiv:2505.13271*, 2025a.

638 Lei Sheng and Shuai-Shuai Xu. Csc-sql: Corrective self-consistency in text-to-sql via reinforcement
 639 learning, 2025b. URL <https://arxiv.org/abs/2505.13271>.
 640

641 Lei Sheng and Shuai-Shuai Xu. Slm-sql: An exploration of small language models for text-to-sql.
 642 *arXiv preprint arXiv:2507.22478*, 2025c.

643 Liang Shi, Zhengju Tang, Nan Zhang, Xiaotong Zhang, and Zhi Yang. A survey on employing large
 644 language models for text-to-sql tasks. *ACM Computing Surveys*, 2024.

645 Vladislav Shkapenyuk, Divesh Srivastava, Theodore Johnson, and Parisa Ghane. Automatic metadata
 646 extraction for text-to-sql, 2025. URL <https://arxiv.org/abs/2505.19988>.
 647

648 Sunanda Somwase, Parismita Das, and S Sudarshan. Data generation for testing complex queries.
 649 *arXiv preprint arXiv:2409.18821*, 2024.
 650

651 Splunk. Ai assistant in observability cloud, September 2025. URL https://www.splunk.com/en_us/products/splunk-ai-assistant-in-observability-cloud.html.
 652

653 Shayan Talaei, Mohammadreza Pourreza, Yu-Chen Chang, Azalia Mirhoseini, and Amin Saberi.
 654 *Chess: Contextual harnessing for efficient sql synthesis*. *arXiv preprint arXiv:2405.16755*, 2024a.
 655

656 Shayan Talaei, Mohammadreza Pourreza, Yu-Chen Chang, Azalia Mirhoseini, and Amin Saberi.
 657 *Chess: Contextual harnessing for efficient sql synthesis*, 2024b. URL <https://arxiv.org/abs/2405.16755>.
 658

659 TiDBCloud. Get Started with Chat2Query API. <https://docs.pingcap.com/tidbcloud/use-chat2query-api>, 2020. Accessed: May 2024.
 660

661 Margus Veanes, Nikolai Tillmann, and Jonathan De Halleux. Qex: Symbolic sql query explorer. In
 662 *International Conference on Logic for Programming Artificial Intelligence and Reasoning*, pp.
 663 425–446. Springer, 2010.
 664

665 Shuxian Wang, Sicheng Pan, and Alvin Cheung. Qed: A powerful query equivalence decider for sql.
 666 *Proceedings of the VLDB Endowment*, 17(11):3602–3614, 2024.
 667

668 Niklas Wretblad, Fredrik Gordh Riseby, Rahul Biswas, Amin Ahmadi, and Oskar Holmström.
 669 Understanding the effects of noise in text-to-sql: An examination of the bird-bench benchmark,
 670 2024.

671 Bohan Zhai, Canwen Xu, Yuxiong He, and Zhewei Yao. Excot: Optimizing reasoning for text-to-sql
 672 with execution feedback, 2025. URL <https://arxiv.org/abs/2503.19988>.
 673

674 Bin Zhang, Yuxiao Ye, Guoqing Du, Xiaoru Hu, Zhishuai Li, Sun Yang, Chi Harold Liu, Rui Zhao,
 675 Ziyue Li, and Hangyu Mao. Benchmarking the text-to-sql capability of large language models: A
 676 comprehensive evaluation, 2024. URL <https://arxiv.org/abs/2403.02951>.
 677

678 Ruiqi Zhong, Tao Yu, and Dan Klein. Semantic evaluation for text-to-sql with distilled test suites.
 679 In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*
 (EMNLP), pp. 396–411, 2020.
 680

681 Aojun Zhou, Ke Wang, Zimu Lu, Weikang Shi, Sichun Luo, Zipeng Qin, Shaoqing Lu, Anya Jia,
 682 Linqi Song, Mingjie Zhan, and Hongsheng Li. Solving challenging math word problems using
 683 GPT-4 code interpreter with code-based self-verification. In *The Twelfth International Conference*
 684 *on Learning Representations*, 2024. URL <https://openreview.net/forum?id=c8McWs4Av0>.
 685

686 Qi Zhou, Joy Arulraj, Shamkant B Navathe, William Harris, and Jinpeng Wu. Spes: A symbolic
 687 approach to proving query equivalence under bag semantics. In *2022 IEEE 38th International*
 688 *Conference on Data Engineering (ICDE)*, pp. 2735–2748. IEEE, 2022.
 689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A RELATED WORK**

703

704 A large number of Text-to-SQL frameworks have been proposed over the last few years by research
 705 groups in academia and industry (Liu et al., 2023; Dong et al., 2023; Chang & Fosler-Lussier, 2023;
 706 TiDBCloud, 2020; Talaei et al., 2024b; Gao et al., 2025; Sequeda et al., 2023; Sheng & Xu, 2025b;
 707 Liu et al., 2025; Shkapenyuk et al., 2025; Zhai et al., 2025). However, evaluation frameworks have
 708 received much less attention. There are two main publicly available platforms: BIRD-SQL (Li et al.,
 709 2024) and Spider (Lei et al., 2024) that are commonly used to evaluate the performance of Text-to-
 710 SQL methods. Their evaluation procedure is performed on predefined database instances, whereas
 711 SPOTIT searches for a separation database instance. A number of evaluation metrics were proposed
 712 to take into account partially correct generated queries (Pinna et al., 2025) or the efficiency of query
 713 executions (Zhang et al., 2024). Recently, (Li et al., 2023a; BIRD, 2025) proposed an iterative
 714 evaluation framework in which the system can interact with the user by asking additional questions
 715 (e.g., to resolve ambiguity). However, the final evaluation of the correctness of the generated SQL
 716 query is still performed on a static database.

717 There are two lines of work in formal equivalence checking for SQL queries: full-fledged and
 718 bounded verification. The full-fledged methods (Chu et al., 2017c; 2018; Zhou et al., 2022; 2024;
 719 Wang et al., 2024) encode queries into specific representations (e.g., algebraic expressions (Chu
 720 et al., 2018; Wang et al., 2024)) and determine equivalence by proving the equivalence of these
 721 representations, thereby guaranteeing equivalence of queries for any possible database. However,
 722 such methods typically support only a limited subset of SQL and cannot generate counterexamples
 723 for non-equivalent queries. In contrast, the bounded verification approaches (Veanes et al., 2010; Chu
 724 et al., 2017a;b; He et al., 2024) check equivalence within a finite search space, making them capable
 725 of handling larger subsets of SQL and identifying counterexamples. To the best of our knowledge,
 726 VERIEQL supports the most expressive SQL fragments and rich integrity constraints, while also
 727 offering extensibility for new features (He et al., 2024). In this work, we significantly extend the
 728 VERIEQL framework to support date and string types as well as a number of common operators for
 729 the Text-to-SQL evaluation task.

730 **Test data generation methods can also be useful for detecting query non-equivalence (Chandra et al.,**
 731 **2015; Somwase et al., 2024; Zhong et al., 2020). However, when the counterexamples require very**
 732 **specific structures (which is the case for many query pairs that passed EX-test but failed SpotIt as**
 733 **seen in App. D), random fuzzing/testing can become unreliable in refuting equivalence. In contrast,**
 734 **SPOTIT systematically searches over the space of possible differentiating databases, finds minimal**
 735 **counter-examples, and provides a formal guarantee: if two SQL queries are considered equivalent,**
 736 **then no counterexample of size less than a fixed number exists. Computational resources permitted,**
 737 **one could in principle run a portfolio of test-based and verification-based equivalence-checking**
 738 **methods in parallel to more quickly detect non-equivalence. This is an orthogonal but interesting**
 739 **future direction.**

756 **B ANALYSIS OF RUNTIME**
757758 Tables 5, 6, 7, 8, and 9 show the effect of different parameters on runtime, including the number of
759 columns, integrity constraints, tables in the databases, the number of sub-queries in the gold SQL,
760 and the number of nodes in the abstract syntax tree in the gold SQL. We found that all parameters
761 except for the number of tables are positively correlated with median runtime.
762

| #columns | Median runtime (s) |
|----------|--------------------|
| 11 | 0.4310 |
| 21 | 0.1701 |
| 31 | 0.1963 |
| 48 | 0.3887 |
| 55 | 0.5576 |
| 64 | 0.3007 |
| 71 | 0.5365 |
| 89 | 0.2672 |
| 94 | 0.3154 |
| 115 | 0.6537 |
| 199 | 0.8006 |

775 Table 5: Median runtime by number of columns
776

| #constraints | Median runtime (s) |
|--------------|--------------------|
| 5 | 0.2842 |
| 7 | 0.1701 |
| 10 | 0.4324 |
| 16 | 0.5576 |
| 17 | 0.3887 |
| 19 | 0.1963 |
| 21 | 0.5365 |
| 36 | 0.6877 |

788 Table 6: Median runtime by number of constraints
789

| #tables | Median runtime (s) |
|---------|--------------------|
| 3 | 0.2842 |
| 4 | 0.4310 |
| 5 | 0.1701 |
| 6 | 0.6537 |
| 7 | 0.8006 |
| 8 | 0.4249 |
| 10 | 0.1963 |
| 13 | 0.3154 |

800 Table 7: Median runtime by number of tables
801

| #subqueries | Median runtime (s) |
|-------------|--------------------|
| 0 | 0.3676 |
| 1 | 0.3972 |
| 2 | 1.9128 |

808 Table 8: Median runtime by number of subqueries
809

| | #AST nodes | Median runtime (s) |
|-----|------------|--------------------|
| 810 | 0–19 | 0.1812 |
| 811 | 20–39 | 0.3408 |
| 812 | 40–59 | 0.3426 |
| 813 | 60–79 | 0.4216 |
| 814 | 80–99 | 0.9683 |
| 815 | 100–119 | 1.1865 |
| 816 | 120–139 | 0.2053 |
| 817 | 140–159 | 1.2395 |
| 818 | | |
| 819 | | |

Table 9: Median runtime by number of AST nodes (buckets of 20)

C LIMITATIONS OF EXISTING BOUNDED SQL EQUIVALENCE CHECKER

While SPOTIT builds on top of and extends VERIEQL, a state-of-the-art bounded verifier that claims to cover the largest SQL fragment, we find that there are still SQL operators which it either does not support or cannot precisely capture. In this section, we describe the features that appear frequently in failure cases of SPOTIT.

- The gold query for question 726 in the BIRD-dev benchmark.

```
820   SELECT superhero.name, height_cm,
821       RANK() OVER (ORDER BY height_cm DESC) AS HeightRank
822   FROM superhero INNER JOIN publisher
823       ON superhero.publisher_id = publisher.id
824   WHERE publisher.publisher_name = "Marvel Comics"
```

SPOTIT does not support the window and analytic functions such as RANK and LAG.

- A SQL query generated by OMNISQL.

```
825   WITH RECURSIVE TimeSeries AS (
826       SELECT '2016-01-01' AS mth
827       UNION ALL
828       SELECT DATE(mth, '+1 month') AS mth FROM TimeSeries
829       WHERE mth < '2017-12-01'
830   ),
831   ...
832   SELECT product_name FROM SalesRatio ... ORDER BY product_name
```

SPOTIT cannot encode recursive common table expressions above.

- Imprecisely encoding for ORDER BY and LIMIT:

Since VERIEQL establishes tables under bag semantics (namely, multi-set), database instances are considered equivalent if they has the same tuples with the same multiplicities. For instance, T_1 and T_2 in Tables 10–11 are equivalent under bag semantics.

However, when VERIEQL symbolically execute ORDER BY A, VERIEQL automatically converts semantics from bag to list in which the order of tuples matters. In such a case, database instances is equivalent iff they are tuple-wise the same. For instance, T_1 and T_2 in Tables 10–11 are not equivalent under list semantics. Furthermore, if ORDER BY A is followed by LIMIT 1, then execution results on T_1 and T_2 are, respectively, R_1 and R_2 .

Table 10: T_1

| | A | B |
|-------|---|---|
| R_1 | 1 | 2 |
| R_2 | 1 | 3 |

Table 11: T_2

| | A | B |
|-------|---|---|
| R_2 | 1 | 3 |
| R_1 | 1 | 2 |

More concretely, consider the gold query Q_1 of question 653 in the BIRD-dev benchmark, a query Q_2 generated by ALPHA, and a counterexample database found by VERIEQL as follows:

864

The gold query Q_1 :

865

```
866     SELECT DisplayName FROM users WHERE id = (
867         SELECT OwnerUserId FROM posts ORDER BY ViewCount DESC LIMIT 1
868     )
```

869

The generated query Q_2 :

870

```
871     SELECT u.displayname AS ownerdisplayname
872         FROM posts AS p INNER JOIN users AS u ON p.owneruserid = u.id
873             ORDER BY p.viewcount DESC LIMIT 1
```

874

Table 12: *posts*

875

| viewcount | owneruserid |
|-----------|-------------|
| Null | 1 |
| Null | 0 |

876

VERIEQL's execution results of Q_1 and Q_2 are shown in Tables 14–15.

877

878

Table 14: VERIEQL's result of Q_1

879

| DisplayName |
|-------------|
| 'A' |

880

881

Table 15: VERIEQL's result of Q_2

882

| ownerdisplayname |
|------------------|
| 'B' |

883

SQLite's execution results of Q_1 and Q_2 are shown in Tables 16–17.

884

885

Table 16: SQLite's result of Q_1

886

887

Table 17: SQLite's result of Q_2

888

| DisplayName |
|-------------|
| 'A' |

889

| ownerdisplayname |
|------------------|
| 'A' |

890

891

892

893

894

Naturally, VERIEQL identifies a spurious counterexample where Q_1 's result is 'B' instead of 'A'. This is because the intermediate table from the FROM clause of Q_2 is shown as Table 18 where the values in column "viewcount" are all Null values. While executing the ORDER BY, VERIEQL does not reorder the tuples of this intermediate table but SQLite engine will swap these two tuples. Therefore, VERIEQL failed in this verification task.

895

896

Table 18: VERIEQL's intermediate table of Q_2

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

| u.DisplayName | p.viewcount |
|---------------|-------------|
| 'A' | Null |
| 'B' | Null |

972 D.3 ADDITIONAL EXAMPLES
973

974 **Example D.1.** Consider the question N_3 and the corresponding SQL queries (Figure 8). The
975 differentiating database found by SPOTIT is shown in Tables 23,24, 25. Note that there is a typo in
976 the evidence. According to external medical sources, the normal range of uric acid levels in females
977 should be defined as less than or equal to 6.50, not greater than. The annotator overlooked this typo,
978 and as a result, the gold SQL is clearly incorrect. \square

979

980 N_3 : "What is the anti Cardiolipin antibody concentration of the female patient
981 with the highest uric acid level in the normal range?"
982 Evidence: "Anti Cardiolipin antibody concentration refers to 'aCL IgG', 'aCL IgM', 'aCL IgA';
983 female patient refers to Sex = F'; highest uric acid level in the normal range refers to $\text{MAX}(\text{UA} > 6.50)$;"
984 /*Gold SQL Q*/:
985

```
SELECT T3.acl_igg, T3.acl_igm, T3.acl_iga
  FROM patient AS T1
  INNER JOIN laboratory AS T2 ON T1.id = T2.id
  INNER JOIN examination AS T3 ON T3.id = T2.id
  WHERE T1.sex = 'F' AND T2.ua > 6.5
  ORDER BY T2.ua DESC
  LIMIT 1
```


986 /*Generated SQL P*/:
987

```
SELECT examination.acl_igg, examination.acl_igm, examination.acl_iga
  FROM patient
  INNER JOIN laboratory ON patient.id = laboratory.id
  INNER JOIN examination ON patient.id = examination.id
  WHERE patient.sex = 'F' AND laboratory.ua <= 6.5
  ORDER BY laboratory.ua DESC
  LIMIT 1
```

995

Figure 8: An example of a query with an incorrect gold SQL.

996

997

998 Table 23: patient (skipped irrelevant columns)

999

1000

1001

| | id | sex | ... |
|---|----|-----|-----|
| 0 | | 'F' | ... |

1002

1003

Table 24: laboratory (skipped irrelevant columns)

1004

1005

1006

| | id | ua | ... |
|---|----|-----|-----|
| 0 | | 6.5 | ... |

1007

1008

Table 25: examination (skipped irrelevant columns)

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

| | id | acl_igg | acl_igm | acl_iga | ... |
|---|----|---------|---------|---------|-----|
| 0 | | 1 | 1 | 1 | ... |

1026 **Example D.2.** Consider the question N_4 and the corresponding SQL queries (Figure 9). The
 1027 differentiating database found by SPOTIT is shown in Tables 26,27. The natural language question
 1028 asks for transactions after January 1st, 2012, which requires excluding January 1st, 2012. However,
 1029 the gold SQL uses a greater-than-or-equal-to condition, which includes 2012/01/01, thus being
 1030 incorrect. \square

1031

1032 N_4 : "Among the transactions made in gas stations in the Czech Republic, how many took place after 2012/1/1?"
 1033 Evidence: "Country code for Czech Republic is 'CZE'."

```
1034     /*Gold SQL Q*/:  

1035     SELECT COUNT(T1.transactionid)  

1036     FROM transactions_1k AS T1  

1037     INNER JOIN gasstations AS T2 ON T1.gasstationid = T2.gasstationid  

1038     WHERE T2.country = 'CZE' AND STRFTIME('%Y', T1.date) ≥ '2012' ;  

1039     /*Generated SQL P*/:  

1040     SELECT COUNT(*)  

1041     FROM transactions_1k AS T  

1042     INNER JOIN gasstations AS G ON T.gasstationid = G.gasstationid  

1043     WHERE G.country = 'CZE' AND T.date > '2012-01-01' ;
```

1042

Figure 9: An example of a query with an incorrect gold SQL.

1043

1044

Table 26: transactions_1k (skipped irrelevant columns)

1045

| transaction_id | gasstation_id | date | ... |
|----------------|---------------|--------------|-----|
| 0 | 0 | '2012-01-01' | ... |

1046

1047

1048

Table 27: gasstations (skipped irrelevant columns)

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080 **Example D.3.** Consider the question N_5 and the corresponding SQL queries (Figure 10). The
 1081 differentiating database found by SPOTIT is shown in Tables 28,29. This example demonstrates an
 1082 incorrect gold SQL, which orders by the latest time (DESC) rather than the earlier time (ASC). This
 1083 directly contradicts the natural language question. \square
 1084

1085 *N_5 : "Which country's gas station had the first paid customer in 2012/8/25?"*
 1086 *Evidence: "2012/8/25" can be represented by '2012-08-25'.*
 1087 /*Gold SQL Q */:
 1088 SELECT T2.country
 1089 FROM transactions_1k AS T1
 1090 INNER JOIN gasstations AS T2 ON T1.gasstationid = T2.gasstationid
 1091 WHERE T1.date = '2012-08-25'
 1092 ORDER BY T1.time DESC
 1093 LIMIT 1;
 1094 /*Generated SQL P */:
 1095 SELECT G.country
 1096 FROM gasstations AS G
 1097 JOIN (
 1098 SELECT gasstationid
 1099 FROM transactions_1k
 1100 WHERE date = '2012-08-25'
 1101 ORDER BY time ASC LIMIT 1
 1102) AS T
 1103 ON G.gasstationid = T.gasstationid;

Figure 10: An example of a query with an incorrect gold SQL.

Table 28: transactions_1k (skipped irrelevant columns)

| | gasstation_id | date | time | ... |
|---|---------------|--------------|------|-----|
| 0 | | '2012-08-25' | 1 | ... |
| 0 | | '2012-08-25' | 2 | ... |

Table 29: gasstations (skipped irrelevant columns)

| | gasstation_id | country | ... |
|---|---------------|---------|-----|
| 0 | | '1' | ... |

1134 **Example D.4.** Consider the question N_6 and the corresponding SQL queries (Figure 11). The
 1135 differentiating database found by SPOTIT is shown in Tables 30, 31. The gold SQL incorrectly
 1136 encodes the exclusive inequality specified in the natural language question by using the BETWEEN
 1137 operator, which leads to inclusive bounds. Thus, the gold SQL is incorrect as it includes values
 1138 outside of the specified range. \square

1139

1140 N_6 : "Please list a patient's platelet level if it is within the normal range
 1141 and if he or she is diagnosed with MCTD"
 1142 Evidence: "PLT > 100 and PLT < 400 means platelet level is within the normal range;
 1143 PLT < 100 and PLT > 400 means platelet level is not within the normal range;
 1144 diagnosed with MCTD refers to Diagnosis = 'MCTD'";

1144 /*Gold SQL Q */:
 1145 SELECT T2.plt
 1146 FROM patient AS T1
 1147 INNER JOIN laboratory AS T2 ON T1.id = T2.id
 1148 WHERE T1.diagnosis = 'MCTD' AND T2.plt BETWEEN 100 AND 400

1144 /*Generated SQL P */:
 1145 SELECT L.plt
 1146 FROM LABORATORY L
 1147 INNER JOIN PATIENT P ON L.id = P.id
 1148 WHERE P.diagnosis = 'MCTD' AND L.plt > 100 AND L.plt < 400

1152

Figure 11: An example of a query with an incorrect gold SQL.

1153

1154

Table 30: patient (skipped irrelevant columns)

1155

1156

| id | diagnosis | ... |
|----|-----------|-----|
| 0 | 'MCTD' | ... |

1158

1159

Table 31: laboratory (skipped irrelevant columns)

1160

1161

| id | plt | ... |
|----|-----|-----|
| 0 | 100 | ... |

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187

1188
 1189 **Example D.5.** Consider the question N_7 and the corresponding SQL queries (Figure 12). The
 1190 differentiating database found by SPOTIT is shown in Tables 32, 33. In this example, the generated
 1191 SQL is incorrect as it is clearly missing the `link_to_major` constraint, filtering only by name. \square
 1192

1193 N_7 : "Please indicate the college of the person whose first name is Katy
 1194 with the link to the major 'rec1N0upiVLy5esTO' "
 1195
 1196 /*Gold SQL Q*/:
 1197 SELECT T2.college
 1198 FROM member AS T1
 1199 INNER JOIN major AS T2 ON T2.major_id = T1.link_to_major
 1200 WHERE T1.link_to_major = 'rec1N0upiVLy5esTO' AND T1.first_name = 'Katy'
 1201 /*Generated SQL P*/:
 1202 SELECT major.college
 1203 FROM member
 1204 INNER JOIN MAJOR ON member.link_to_major = major.major_id
 1205 WHERE member.first_name = 'Katy'

Figure 12: An example of a query with an incorrect generated SQL.

Table 32: member (skipped irrelevant columns)

| link_to_major | first_name | ... |
|---------------|------------|-----|
| '1' | 'Katy' | ... |

Table 33: major (skipped irrelevant columns)

| major.id | college | ... |
|----------|---------|-----|
| 1 | '0' | ... |

1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241

1242 **Example D.6.** Consider the question N_8 and the corresponding SQL queries (Figure 13). The
 1243 differentiating database found by SPOTIT is shown in Tables 34, 35. In this example, the generated
 1244 SQL only checks whether the patient was diagnosed with SLE on January 1st, 1997. However, the
 1245 natural language question also asks for the patient’s original diagnose at their first hospital visit.
 1246 Since the generated SQL doesn’t include this condition, it’s incorrect as it could return a diagnoses
 1247 from a later visit rather than the patient’s first one. \square

1248

1249 N_8 : “For the patient who was diagnosed SLE on 1997/1/27, what was his/her original diagnose when he/she came to the hospital for the
 1250 first time?”
 1251 Evidence: “‘SLE’ and original diagnose refers to Diagnosis; 1997/1/27 refers to ‘Examination Date’ = ‘1997-01-27’; first came to the
 1251 hospital refers to patient. First Date.”

1252

```
/*Gold SQL Q*/:  

SELECT T1.diagnosis  

FROM patient AS T1  

INNER JOIN examination AS T2 ON T1.id = T2.id  

WHERE T1.id = (  

  SELECT id  

  FROM examination  

  WHERE examination_date = '1997-01-27' AND diagnosis = 'SLE'  

) AND T2.examination_date = T1.first_date;  

/*Generated SQL P*/:  

SELECT T2.diagnosis  

FROM examination AS T1  

INNER JOIN patient AS T2 ON T1.id = T2.id  

WHERE T1.diagnosis = 'SLE' AND T1.examination_date = '1997-01-27';
```

1264

Figure 13: An example of a query with an incorrect generated SQL.

1266

1267

Table 34: patient (skipped irrelevant columns)

1269

1270

| id | diagnosis | first_date | ... |
|----|-----------|--------------|-----|
| 0 | ‘1’ | ‘1997-01-26’ | ... |

1271

1272

Table 35: examination (skipped irrelevant columns)

1273

1274

1275

| id | examination_date | diagnosis | ... |
|----|------------------|-----------|-----|
| 0 | ‘1997-01-27’ | ‘SLE’ | ... |

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

1296 **Example D.7.** Consider the question N_9 and the corresponding SQL queries (Figure 14). The
 1297 differentiating database found by SPOTIT is shown in Tables 36,37. This is an example of an
 1298 ambiguous question. The term 'members' can be interpreted in at least two ways: any student
 1299 who is a part of the club, or more specifically, students in the club with the recorded position of
 1300 'member'. While the gold SQL takes the second interpretation, filtering on $T2.position = 'Member'$,
 1301 it's just as reasonable to assume that all students in the club are members, and leave out a secondary
 1302 filter. Coupled with the lack of evidence, the resulting difference in queries is most likely due to the
 1303 ambiguity of the natural language question. Hence, it's been marked as an ambiguous question. \square

1304

1305 N_9 : "List the last name of members with a major in environmental engineering
 1306 and include its department and college name.
 1307 Evidence: "Environmental Engineering' is the major name"
 1308 /*Gold SQL Q*/:
 1309 SELECT T2.last_name, T1.department, T1.college
 1310 FROM major AS T1
 1311 INNER JOIN member AS T2 ON T1.major_id = T2.link_to_major
 1312 WHERE T2.position = 'Member' AND T1.major.name = 'Enviornmental Engineering'
 1313 /*Generated SQL P*/:
 1314 SELECT T1.last_name, T2.department, T2.college
 1315 FROM member AS T1
 1316 INNER JOIN major AS T2 ON T1.link_to_major = T2.major_id
 1317 WHERE T2.major.name = 'Enviornmental Engineering'

Figure 14: An example of an ambiguous question.

1318

Table 36: major (skipped irrelevant columns)

1319

| major_id | major_name | department | college | ... |
|----------|-----------------------------|------------|---------|-----|
| 0 | 'Environmental Engineering' | '1' | '1' | ... |

1322

Table 37: member (skipped irrelevant columns)

1323

| last_name | link_to_major | position | ... |
|-----------|---------------|----------|-----|
| '1' | 0 | '1' | ... |

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350
 1351 **Example D.8.** Consider the question N_{10} and the corresponding SQL queries (Figure 15). The
 1352 differentiating database found by SPOTIT is shown in Tables 38, 39. This example is marked as
 1353 ambiguous because the natural language question is underspecified. If the intent is to return the
 1354 legal status of every valid artifact card, which is a reasonable interpretation, than the generated
 1355 SQL would be correct. However, if the intent is to return the set of unique legal statuses across valid
 1356 artifact cards, than the gold SQL is correct. \square

1357 N_{10} : "For artifact type of cards that do not have multiple faces on the same card, state its legalities status for vintage play format."
 1358 Evidence: "Artifact type of cards refers to types = 'Artifact'; card does not have multiple faces on the same card refers to side is NULL";
 1359 vintage play format refers to format = 'vintage';"

1360 /*Gold SQL Q */:
 1361 **SELECT** T2.status
 1362 **FROM** cards **AS** T1
 1363 **INNER JOIN** legalities **AS** T2 **ON** T1.uuid = T2.uuid
 1364 **WHERE** T1.type = 'Artifact' **AND** T2.format = 'vintage' **AND** T1.side IS **NULL**;
 1365 /*Generated SQL P */:
 1366 **SELECT** T2.status
 1367 **FROM** cards **AS** T1
 1368 **JOIN** legalities **AS** T2 **ON** T1.uuid = T2.uuid
 1369 **WHERE** T1.type = 'Artifact' **AND** T1.side IS **NULL** **AND** T2.format = 'vintage';

Figure 15: An example of an ambiguous question.

Table 38: cards (skipped irrelevant columns)

| uuid | type | side | ... |
|------|------------|------|-----|
| '0' | 'Artifact' | NULL | ... |

Table 39: legalities (skipped irrelevant columns)

| uuid | format | status | ... |
|------|-----------|--------|-----|
| '0' | 'vintage' | '1' | ... |
| '0' | 'vintage' | '1' | ... |

1404
 1405 **Example D.9.** Consider the question N_{11} and the corresponding SQL queries (Figure 16). The
 1406 differentiating database found by SPOTIT is shown in Tables 40, 41. This example is considered
 1407 ambiguous because the natural language question and evidence do not specify a tie-breaking rule. In
 1408 the case that there are two comments on valid posts with a tied high score, a query with `LIMIT 1` may
 1409 return either comment. This is why the generated and gold SQL return different results. Since the
 1410 difference arises solely from a lack of specificity, this example is marked as ambiguous. \square

1411 N_{11} : "Among the posts with views ranging from 100 to 150, what is the comment with the highest score?"
 1412 Evidence: "Views ranging from 100 to 150 refers to `ViewCount BETWEEN 100 and 150`; comment with the highest score refers to `Text`
 1413 where `MAX(Score)`;"
 1414 /*Gold SQL Q^* :
 1415 `SELECT text`
 1416 `FROM comments`
 1417 `WHERE postId IN (`
 1418 `SELECT id`
 1419 `FROM posts`
 1420 `WHERE viewCount BETWEEN 100 AND 150`
 1421 `) ORDER BY score DESC`
 1422 `LIMIT 1`
 1423 /*Generated SQL P^* :
 1424 `SELECT T2.text`
 1425 `FROM posts AS T1`
 1426 `INNER JOIN comments AS T2 ON T1.id = T2.postId`
 1427 `WHERE T1.viewCount BETWEEN 100 AND 150`
 1428 `ORDER BY T2.score DESC`
 1429 `LIMIT 1`

Figure 16: An example of an ambiguous question.

Table 40: comments (skipped irrelevant columns)

| postId | score | text | ... |
|--------|-------|------|-----|
| 0 | 1 | '1' | ... |
| 1 | 1 | '2' | ... |

Table 41: posts (skipped irrelevant columns)

| id | viewCount | ... |
|----|-----------|-----|
| 0 | 100 | ... |
| 1 | 100 | ... |

1458 E SEMANTICS
1459

| 1461 $\llbracket E \rrbracket :: \text{Database } D \rightarrow \text{Relation} \rightarrow \text{Value}$ | |
|---|--|
| 1463 $\llbracket \text{ToInt}(E) \rrbracket_{D, xs}$ | = let $v = \llbracket E \rrbracket_{D, xs}$ in 1464 $\text{ite}(v = \text{Null} \vee \text{IsInt}(v), v,$ 1465 $\text{ite}(\text{IsStr}(v), \llbracket \text{StrToInt}(v) \rrbracket_{D, xs}, \llbracket \text{DateToInt}(v) \rrbracket_{D, xs}))$ |
| 1466 $\llbracket \text{ToDate}(E) \rrbracket_{D, xs}$ | = let $v = \llbracket E \rrbracket_{D, xs}$ in 1467 $\text{ite}(v = \text{Null} \vee \text{IsDate}(v), v,$ 1468 $\text{ite}(\text{IsInt}(v), \llbracket \text{IntToDate}(v) \rrbracket_{D, xs}, \llbracket \text{StrToDate}(v) \rrbracket_{D, xs}))$ |
| 1469 $\llbracket \text{ToStr}(E) \rrbracket_{D, xs}$ | = let $v = \llbracket E \rrbracket_{D, xs}$ in 1470 $\text{ite}(v = \text{Null} \vee \text{IsStr}(v), v,$ 1471 $\text{ite}(\text{IsInt}(v), \llbracket \text{IntToStr}(v) \rrbracket_{D, xs}, \llbracket \text{DateToStr}(v) \rrbracket_{D, xs}))$ |
| 1472 $\llbracket \text{DateToInt}(vs) \rrbracket_{D, xs}$ | = $\text{ite}(vs = \text{Null}, \text{Null}, vs[0] * 10^4 + vs[1] * 10^2 + vs[2])$ |
| 1473 $\llbracket \text{StrToInt}(s) \rrbracket_{D, xs}$ | = let 1474 $v = \text{ite}(\text{IsDigits}(s), \text{StrToInt}(s),$ 1475 $\text{ite}(s[0] = “-” \wedge \text{IsDigits}(s[1:]), -\text{StrToInt}(s), 0))$ 1476 in $\text{ite}(s = \text{Null}, \text{Null}, v)$ |
| 1477 $\llbracket \text{IntToStr}(v) \rrbracket_{D, xs}$ | = $\text{ite}(v = \text{Null}, \text{Null}, \text{IntToStr}(v))$ |
| 1478 $\llbracket \text{DateToStr}(vs) \rrbracket_{D, xs}$ | = let 1479 $y = \text{IntToStr}(vs[0]),$ 1480 $m = \text{ite}(vs[1] \leq 9, “0” + \text{IntToStr}(vs[1]), \text{IntToStr}(vs[1])),$ 1481 $d = \text{ite}(vs[2] \leq 9, “0” + \text{IntToStr}(vs[2]), \text{IntToStr}(vs[2]))$ 1482 in $\text{ite}(vs = \text{Null}, \text{Null}, y + “-” + m + “-” + d)$ |
| 1483 $\llbracket \text{IntToDate}(v) \rrbracket_{D, xs}$ | = let $v_1 = \lfloor v/10^4 \rfloor, v_2 = \lfloor (v \% 10^4)/10^2 \rfloor, v_3 = v \% 10^2$ in 1484 $\text{ite}(v = \text{Null} \vee \text{IsValidDate}(v), \text{Null}, [v_1, v_2, v_3])$ |
| 1485 $\llbracket \text{StrToDate}(s) \rrbracket_{D, xs}$ | = let $v = \llbracket \text{StrToInt}(s) \rrbracket_{D, xs}$ in 1486 $\text{ite}(s = \text{Null}, \text{Null}, \llbracket \text{IntToDate}(v) \rrbracket_{D, xs})$ |
| 1487 $\llbracket E_1 \diamond E_2 \rrbracket_{D, xs}$ | = let 1488 $v_1 = \llbracket \text{ToInt}(E_1) \rrbracket_{D, xs}$ and $v_2 = \llbracket \text{ToInt}(E_2) \rrbracket_{D, xs},$ 1489 in $\text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, v_1 \diamond v_2)$ |
| 1490 $\llbracket \text{SubStr}(E_1, E_2, E_3) \rrbracket_{D, xs}$ | = let 1491 $e_i = \llbracket E_i \rrbracket_{D, xs}, e'_1 = \llbracket \text{ToStr}(e_1) \rrbracket_{D, xs}, l = \text{len}(e'_1),$ 1492 $e'_2 = \llbracket \text{ToInt}(e_2) \rrbracket_{D, xs}, e'_3 = \llbracket \text{ToInt}(e_3) \rrbracket_{D, xs},$ 1493 $v = \text{ite}(-l \leq e'_2 < 0, e_2 + l, \text{ite}(0 < e'_2 \leq l, e'_2 - 1, l + 1)),$ 1494 $s = \text{ite}(v = 0 \vee v < -l \vee v > l \vee e'_3 \leq 0, \text{Null},$ 1495 $\text{ite}(e'_3 \geq l - v, e'_1[v:l], e'_1[v:2v + e'_3]))$ 1496 in $\text{ite}(e_1 = \text{Null} \vee \text{IsStr}(e_2) \vee \text{IsStr}(e_3), \text{Null}, s)$ |
| 1497 $\llbracket \text{Concat}(E_1, E_2) \rrbracket_{D, xs}$ | = let $v_1 = \llbracket \text{ToStr}(E_1) \rrbracket_{D, xs}$ and $v_2 = \llbracket \text{ToStr}(E_2) \rrbracket_{D, xs}$ in 1498 $\text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, \text{Concat}(v_1, v_2))$ |
| 1499 $\llbracket \text{Strftime}(\kappa, E) \rrbracket_{D, xs}$ | = let $v = \llbracket \text{ToDate}(E) \rrbracket_{D, xs}$ in 1500 $\text{ite}(\kappa = “%Y”, v[0], \text{ite}(\kappa = “%M”, v[1], v[2]))$ |
| 1501 $\llbracket \text{JulianDay}(E) \rrbracket_{D, xs}$ | = let $v = \llbracket E \rrbracket_{D, xs}$ in $\text{ToJulianDay}(v)$, if $\text{IsDate}(v)$ |
| 1502 $\llbracket \text{DateShift}(E, i, \delta) \rrbracket_{D, xs}$ | = let $v = \llbracket E \rrbracket_{D, xs}$ in $\text{DateAdd}(v, i, \delta)$, if $\text{IsDate}(v)$ |

| 1500 $\llbracket \phi \rrbracket :: \text{Database } D \rightarrow \text{Relation} \rightarrow \text{Bool} \cup \text{Null}$ | |
|--|--|
| 1501 $\llbracket \text{PrefixOf}(s, E) \rrbracket_{D, xs}$ | = let $v = \llbracket \text{ToStr}(E) \rrbracket_{D, xs}$ in $\text{ite}(v = \text{Null}, \text{Null}, \text{PrefixOf}(s, v))$ |
| 1502 $\llbracket \text{SuffixOf}(s, E) \rrbracket_{D, xs}$ | = let $v = \llbracket \text{ToStr}(E) \rrbracket_{D, xs}$ in $\text{ite}(v = \text{Null}, \text{Null}, \text{SuffixOf}(s, v))$ |
| 1503 $\llbracket \text{Like}(s, E) \rrbracket_{D, xs}$ | = let $v = \llbracket \text{ToStr}(E) \rrbracket_{D, xs}$ in $\text{ite}(v = \text{Null}, \text{Null}, \text{RegexMatch}(s, v))$ |
| 1504 $\llbracket \text{Contain}(s, E) \rrbracket_{D, xs}$ | = let $s' = \text{Concat}(“%”, s, “%”)$ in $\llbracket \text{Like}(s', E) \rrbracket_{D, xs}$ |
| 1505 $\llbracket E_1 \odot E_2 \rrbracket_{D, xs}$ | = let $v_1 = \llbracket E_1 \rrbracket_{D, xs}$ and $v_2 = \llbracket E_2 \rrbracket_{D, xs}$ in 1506 $\text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, v_1 \odot v_2)$, if $\text{Type}(v_1) = \text{Type}(v_2)$ |

1507 Figure 17: Formal semantics for extended expressions and predicates. The `IsValidDate` function
1508 checks whether a string represent a date within the supported date range of a database engine. The
1509 `ToJulianDay` function converts a date to a Julian day and the `DateAdd` function move a date-value by
1510 modifier arguments i and δ . The definition of these two functions are shown in Appendix G.
1511

1512 F ENCODING
1513

| | |
|---|---|
| 1515 $\llbracket \text{ToInt}(E) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket E \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(v = \text{Null} \vee \text{IsInt}(v), v,$ $\quad \quad \text{ite}(\text{IsStr}(v), \llbracket \text{StrToInt}(v) \rrbracket_{S,\Gamma,\tau}, \llbracket \text{DateToInt}(v) \rrbracket_{S,\Gamma,\tau}))$ |
| 1516 $\llbracket \text{ToDate}(E) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket E \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(v = \text{Null} \vee \text{IsDate}(v), v,$ $\quad \quad \text{ite}(\text{IsInt}(v), \llbracket \text{IntToDate}(v) \rrbracket_{S,\Gamma,\tau}, \llbracket \text{StrToDate}(v) \rrbracket_{S,\Gamma,\tau}))$ |
| 1517 $\llbracket \text{ToStr}(E) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket E \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(v = \text{Null} \vee \text{IsStr}(v), v,$ $\quad \quad \text{ite}(\text{IsInt}(v), \llbracket \text{IntToStr}(v) \rrbracket_{S,\Gamma,\tau}, \llbracket \text{DateToStr}(v) \rrbracket_{S,\Gamma,\tau}))$ |
| 1518 $\llbracket \text{DateToInt}(vs) \rrbracket_{S,\Gamma,\tau}$ | $= \text{ite}(vs = \text{Null}, \text{Null}, vs[0] * 10^4 + vs[1] * 10^2 + vs[2])$ |
| 1519 $\llbracket \text{StrToInt}(s) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let}$ $\quad s_1 = s[1:z3.\text{Length}(s)], v_1 = z3.\text{StrToInt}(s_1),$ $\quad v = \text{ite}(s[0] = “-”, -v_1, z3.\text{StrToInt}(s)),$ $\quad \Phi = \text{ite}(v < 0, z3.\text{IntToStr}(-v) = v_1, z3.\text{IntToStr}(v) = s),$ $\quad \text{in } \text{ite}(s = \text{Null}, \text{Null}, \text{ite}(\Phi, v, 0))$ |
| 1520 $\llbracket \text{IntToStr}(v) \rrbracket_{S,\Gamma,\tau}$ | $= \text{ite}(v = \text{Null}, \text{Null}, z3.\text{IntToStr}(v))$ |
| 1521 $\llbracket \text{DateToStr}(vs) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } y = z3.\text{IntToStr}(vs[0]),$ $\quad m = \text{ite}(vs[1] \leq 9, “0” + z3.\text{IntToStr}(vs[1]),$ $\quad \quad z3.\text{IntToStr}(vs[1])),$ $\quad d = \text{ite}(vs[2] \leq 9, “0” + z3.\text{IntToStr}(vs[2]),$ $\quad \quad z3.\text{IntToStr}(vs[2]))$ $\quad \text{in } \text{ite}(vs = \text{Null}, \text{Null}, y + “-” + m + “-” + d)$ |
| 1522 $\llbracket \text{IntToDate}(v) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } y = \text{fdiv}(v, 10^4), m = \text{fdiv}(v \% 10^4, 10^2), d = v \% 10^2,$ $\quad \Phi_0 = y \% 4 = 0 \wedge (y \% 100 \neq 0 \vee y \% 400 = 0)$ $\quad \Phi_1 = \text{MIN_YEAR} \leq y \leq \text{MAX_YEAR},$ $\quad \Phi_2 = 1 \leq m \leq 12,$ $\quad \Phi_3 = 1 \leq d \wedge (\vee_{c \in \{1,3,5,7,8,10,12\}} m = c \rightarrow d \leq 31)$ $\quad \quad \wedge (m = 2 \rightarrow d \leq 28 + \text{ite}(\Phi_0, 1, 0))$ $\quad \quad \wedge (\vee_{c \in \{4,6,9,11\}} m = c \rightarrow d \leq 30)$ $\quad \text{in } \text{ite}(v = \text{Null} \vee \neg(\Phi_1 \wedge \Phi_2 \wedge \Phi_3), \text{Null}, [y, m, d])$ |
| 1523 $\llbracket \text{StrToDate}(s) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket \text{StrToInt}(s) \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(s = \text{Null}, \text{Null}, \llbracket \text{IntToDate}(v) \rrbracket_{S,\Gamma,\tau})$ |
| 1524 $\llbracket E_1 \diamond E_2 \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v_1 = \llbracket \text{ToInt}(E_1) \rrbracket_{S,\Gamma,\tau} \text{ and } v_2 = \llbracket \text{ToInt}(E_2) \rrbracket_{S,\Gamma,\tau},$ $\quad \text{in } \text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, v_1 \diamond v_2)$ |
| 1525 $\llbracket \text{SubStr}(E_1, E_2, E_3) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let}$ $\quad e_i = \llbracket E_i \rrbracket_{S,\Gamma,\tau}, e'_1 = \llbracket \text{ToStr}(e_1) \rrbracket_{S,\Gamma,\tau}, l = z3.\text{Length}(e'_1),$ $\quad e'_2 = \llbracket \text{ToInt}(e_2) \rrbracket_{S,\Gamma,\tau}, e'_3 = \llbracket \text{ToInt}(e_3) \rrbracket_{S,\Gamma,\tau},$ $\quad v = \text{ite}(-l \leq e'_2 < 0, e_2 + l, \text{ite}(0 < e'_2 \leq l, e'_2 - 1, l + 1)),$ $\quad s = \text{ite}(v = 0 \vee v < -l \vee v > l \vee e'_3 \leq 0, \varepsilon,$ $\quad \quad \text{ite}(e'_3 \geq l - v, e'_1[v:l], e'_1[v:2v + e'_3]))$ $\quad \text{in } \text{ite}(e_1 = \text{Null} \vee \text{IsStr}(e_2) \vee \text{IsStr}(e_3), \text{Null}, s)$ |
| 1526 $\llbracket \text{Concat}(E_1, E_2) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v_1 = \llbracket \text{ToStr}(E_1) \rrbracket_{S,\Gamma,\tau} \text{ and } v_2 = \llbracket \text{ToStr}(E_2) \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, z3.\text{Concat}(v_1, v_2))$ |
| 1527 $\llbracket \text{Strftime}(\kappa, E) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket \text{ToDate}(E) \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(v = \text{Null}, \text{Null},$ $\quad \quad \text{ite}(\kappa = “\%Y”, v[0], \text{ite}(\kappa = “\%M”, v[1], v[2])))$ |
| 1528 $\llbracket \text{JulianDay}(E) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket E \rrbracket_{S,\Gamma,\tau}, y = \text{ite}(v[1] \leq 2, v[0] - 1, v[0]),$ $\quad m = \text{ite}(v[1] \leq 2, v[1] + 12, v[1]), d = v[2],$ $\quad c = 2 - \text{fdiv}(y, 100) + \text{fdiv}(y, 400),$ $\quad a_1 = \text{fdiv}(36525 * (y + 4716), 10^2),$ $\quad a_2 = \text{fdiv}(306001 * (m + 1), 10^4),$ $\quad \text{in } a_1 + a_2 + d + c - 1524.5, \text{ if IsDate}(v)$ |
| 1529 $\llbracket \text{DateShift}(E, i, \delta) \rrbracket_{S,\Gamma,\tau}$ | $= \text{let } v = \llbracket E \rrbracket_{S,\Gamma,\tau} \text{ in}$ $\quad \text{ite}(\delta = “\text{Year}”, \text{DateShiftByYears}(v, i),$ $\quad \quad \text{ite}(\delta = “\text{Month}”, \text{DateShiftByMonths}(v, i),$ $\quad \quad \quad \text{DateShiftByDays}(v, i)))$ |

Figure 18: Symbolic encoding for extended expressions. The floor division function is defined as $\text{fdiv}(x, y) = \text{ite}(x \% y = 0, x/y, (x - x \% y)/y)$. For clarity, we overload `IsInt`, `IsStr` and `IsDate` to check whether formulas represent integers, strings and dates, respectively. Type conversions and string manipulations are handled using the built-in functions of Z3.

```

1566  $\llbracket \text{PrefixOf}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{let } v = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} \text{ in } \text{ite}(v = \text{Null}, \text{Null}, \text{z3.PrefixOf}(s, v))$ 
1567  $\llbracket \text{SuffixOf}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{let } v = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} \text{ in } \text{ite}(v = \text{Null}, \text{Null}, \text{z3.SuffixOf}(s, v))$ 
1568  $\llbracket \text{Like}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{let } v = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} \text{ in }$ 
1569  $\qquad \text{ite}(v = \text{Null}, \text{Null}, \text{z3.RegexMatch}(s, v))$ 
1570  $\llbracket \text{Contain}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{let } v = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} \text{ and } s' = \text{Concat}(\text{“.”}, s, \text{“.”}) \text{ in }$ 
1571  $\qquad \text{ite}(v = \text{Null}, \text{Null}, \text{z3.RegexMatch}(s', v))$ 
1572  $\llbracket E_1 \odot E_2 \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{let } v_1 = \llbracket E_1 \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} \text{ and } v_2 = \llbracket E_2 \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} \text{ in }$ 
1573  $\qquad \text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \perp, v_1 \odot v_2), \text{ if Type}(v_1) = \text{Type}(v_2)$ 
1574
1575
1576
1577
1578
1579

```

Figure 19: Symbolic encoding for extended predicates.

G PROOF

In this section, we provide the proof of theorems in the main paper.

Theorem 1 (Correctness of expression encoding). *Let D be a database over schema \mathcal{S} , xs be a tuple list, and E be an expression. Consider a symbolic database Γ over \mathcal{S} , a list of symbolic tuples \mathcal{T} , and E 's symbolic encoding $\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$. For any satisfying interpretation \mathcal{I} with $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs$, evaluating the expression E over the database D and the tuple list xs yields the interpretation of E 's symbolic encoding $\mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$, i.e., $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs \Rightarrow \llbracket E \rrbracket_{D, xs} = \mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$.*

Lemma 1. *Suppose $\llbracket E \rrbracket_{D, xs} = v$, then $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs \Rightarrow \llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$ is true iff $\llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v$ and $\mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) = v$.*

Proof. Theorem 1 is proved by proving Lemma 1. By structural induction on E .

1. Base cases and some inductive cases are proved in He et al. (2024).

2. Inductive case: $E = \text{ToInt}(E)$

$\llbracket \text{ToInt}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null} \vee \text{IsInt}(v), v, \text{ite}(\text{IsStr}(v), \llbracket \text{StrToInt}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}},$
 $\llbracket \text{DateToInt}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}))$ where $v = \llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ by Figure 18. $\llbracket \text{ToInt}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null} \vee \text{IsInt}(v'), v', \text{ite}(\text{IsStr}(v'), \llbracket \text{StrToInt}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})},$
 $\llbracket \text{DateToInt}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}))$ where $v' = \llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By inductive hypothesis, we have $\mathcal{I}(v) = \mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) = \llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'$. Therefore,

$$\begin{aligned}
\mathcal{I}(\llbracket \text{ToInt}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null} \vee \text{IsInt}(v), v, \text{ite}(\text{IsStr}(v), \llbracket \text{StrToInt}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}, \\
&\qquad \llbracket \text{DateToInt}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}))) \\
&= \text{ite}(\mathcal{I}(v) = \text{Null} \vee \mathcal{I}(\text{IsInt}(v)), \mathcal{I}(v), \text{ite}(\mathcal{I}(\text{IsStr}(v)), \\
&\qquad \mathcal{I}(\llbracket \text{StrToInt}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}), \mathcal{I}(\llbracket \text{DateToInt}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}))) \\
&= \text{ite}(\mathcal{I}(v) = \text{Null} \vee \text{IsInt}(\mathcal{I}(v)), \mathcal{I}(v), \text{ite}(\text{IsStr}(\mathcal{I}(v)), \\
&\qquad \llbracket \text{StrToInt}(\mathcal{I}(v)) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}, \llbracket \text{DateToInt}(\mathcal{I}(v)) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})})) \\
&= \text{ite}(v' = \text{Null} \vee \text{IsInt}(v'), v', \text{ite}(\text{IsStr}(v'), \\
&\qquad \llbracket \text{StrToInt}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}, \llbracket \text{DateToInt}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})})) \\
&= \llbracket \text{ToInt}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}
\end{aligned}$$

3. Inductive case: $E = \text{ToDate}(E)$

$\llbracket \text{ToDate}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null} \vee \text{IsDate}(v), v, \text{ite}(\text{IsInt}(v), \llbracket \text{IntToDate}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}},$
 $\llbracket \text{StrToDate}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}))$ where $v = \llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ by Figure 18. $\llbracket \text{ToDate}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null} \vee \text{IsDate}(v'), v', \text{ite}(\text{IsInt}(v'), \llbracket \text{IntToDate}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})},$
 $\llbracket \text{StrToDate}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}))$ where $v' = \llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By inductive hypothesis, we have

1620 $\mathcal{I}(v) = \mathcal{I}(\llbracket E \rrbracket_{S,\Gamma,\mathcal{T}}) = \llbracket E \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} = v'$. Therefore,

1621
$$\begin{aligned} \mathcal{I}(\llbracket \text{ToDate}(E) \rrbracket_{S,\Gamma,\mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null} \vee \text{IsDate}(v), v, \text{ite}(\text{IsInt}(v), \\ &\quad \llbracket \text{IntToDate}(v) \rrbracket_{S,\Gamma,\mathcal{T}}, \llbracket \text{StrToDate}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null} \vee \mathcal{I}(\text{IsDate}(v)), \mathcal{I}(v), \text{ite}(\mathcal{I}(\text{IsInt}(v)), \\ &\quad \mathcal{I}(\llbracket \text{IntToDate}(v) \rrbracket_{S,\Gamma,\mathcal{T}}, \mathcal{I}(\llbracket \text{StrToDate}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null} \vee \text{IsDate}(\mathcal{I}(v)), \mathcal{I}(v), \text{ite}(\text{IsInt}(\mathcal{I}(v)), \\ &\quad \mathcal{I}(\llbracket \text{IntToDate}(v) \rrbracket_{S,\Gamma,\mathcal{T}}, \mathcal{I}(\llbracket \text{StrToDate}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))) \\ &= \text{ite}(v' = \text{Null} \vee \text{IsDate}(v'), v', \text{ite}(\text{IsInt}(v'), \\ &\quad \llbracket \text{IntToDate}(v') \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})}, \llbracket \text{StrToDate}(v') \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})})) \\ &= \llbracket \text{ToDate}(E) \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} \end{aligned}$$

1630

1631 4. Inductive case: $E = \text{ToStr}(E)$

1632 $\llbracket \text{ToStr}(E) \rrbracket_{S,\Gamma,\mathcal{T}} = \text{ite}(v = \text{Null} \vee \text{IsStr}(v), v, \text{ite}(\text{IsInt}(v), \llbracket \text{IntToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}},$

1633 $\llbracket \text{DateToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))$ where $v = \llbracket E \rrbracket_{S,\Gamma,\mathcal{T}}$ by Figure 18.

1634 $\llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null} \vee \text{IsStr}(v'), v', \text{ite}(\text{IsInt}(v'),$

1635 $\llbracket \text{IntToStr}(v') \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})}, \llbracket \text{DateToStr}(v') \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})}))$ where $v' = \llbracket E \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})}$ by Figure 17. By inductive hypothesis, we have

1637 $\mathcal{I}(v) = \mathcal{I}(\llbracket E \rrbracket_{S,\Gamma,\mathcal{T}}) = \llbracket E \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} = v'$. Therefore,

1638
$$\begin{aligned} \mathcal{I}(\llbracket \text{ToStr}(E) \rrbracket_{S,\Gamma,\mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null} \vee \text{IsStr}(v), v, \text{ite}(\text{IsInt}(v), \\ &\quad \llbracket \text{IntToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}, \llbracket \text{DateToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null} \vee \mathcal{I}(\text{IsStr}(v)), \mathcal{I}(v), \text{ite}(\mathcal{I}(\text{IsInt}(v)), \\ &\quad \mathcal{I}(\llbracket \text{IntToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}, \mathcal{I}(\llbracket \text{DateToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null} \vee \text{IsStr}(\mathcal{I}(v)), \mathcal{I}(v), \text{ite}(\text{IsInt}(\mathcal{I}(v)), \\ &\quad \mathcal{I}(\llbracket \text{IntToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}, \mathcal{I}(\llbracket \text{DateToStr}(v) \rrbracket_{S,\Gamma,\mathcal{T}}))) \\ &= \text{ite}(v' = \text{Null} \vee \text{IsStr}(v'), v', \text{ite}(\text{IsInt}(v'), \\ &\quad \llbracket \text{IntToStr}(v') \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})}, \llbracket \text{DateToStr}(v') \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})})) \\ &= \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} \end{aligned}$$

1648 5. Inductive case: $E = \text{DateToInt}(vs)$

1649 $\llbracket \text{DateToInt}(vs) \rrbracket_{S,\Gamma,\mathcal{T}} = \text{ite}(vs = \text{Null}, \text{Null}, vs[0]*10^4 + vs[1]*10^2 + vs[2])$ by Figure 18.

1650 $\llbracket \text{DateToInt}(vs) \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} = \text{ite}(vs = \text{Null}, \text{Null}, vs[0] * 10^4 + vs[1] * 10^2 + vs[2])$ by Figure 17. Therefore, $\mathcal{I}(\llbracket \text{DateToInt}(vs) \rrbracket_{S,\Gamma,\mathcal{T}}) = \text{ite}(vs = \text{Null}, \text{Null}, vs[0] * 10^4 + vs[1] * 10^2 + vs[2]) = \llbracket \text{DateToInt}(vs) \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})}$.

1654 6. Inductive case: $E = \text{StrToInt}(s)$

1655 $\llbracket \text{StrToInt}(s) \rrbracket_{S,\Gamma,\mathcal{T}} = \text{ite}(s = \text{Null}, \text{Null}, \text{ite}(\Phi, v, 0))$ where $s_1 = s[1 : \text{z3.Length}(s)]$,

1656 $v_1 = \text{z3.StrToInt}(s_1)$, $v = \text{ite}(s[0] = “-”, -v_1, \text{z3.StrToInt}(s))$, and $\Phi = \text{ite}(v < 0, \text{z3.IntToStr}(-v) = v_1, \text{z3.IntToStr}(v) = s)$ by Figure 18.

1657 $\llbracket \text{StrToInt}(s) \rrbracket_{\mathcal{I}(\Gamma),\mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null}, \text{Null}, v')$ where $v' = \text{ite}(\text{IsDigits}(s), \text{StrToInt}(s), \text{ite}(s[0] = “-” \wedge \text{IsDigits}(s[1 :]), -\text{StrToInt}(s), 0))$ by Figure 17.

1660 On the one hand, the Z3 builtin function $\text{z3.StrToInt}(s) = \text{StrToInt}(s)$ if $\text{StrToInt}(s) \geq 0$;

1661 otherwise, $\text{z3.StrToInt}(s) = -1$. To show our encoding precisely capture semantics of

1662 *SQL's type conversion from strings to integers*, let us discuss it in three cases:

1663

- (a) If $\text{StrToInt}(s) \geq 0$, then $v = \text{z3.StrToInt}(s) = \text{StrToInt}(s)$ and Φ holds. Thus, $\text{ite}(\Phi, v, 0) = v$.
- (b) If $\text{StrToInt}(s) < 0$, then $v = -v_1$ and $\Phi = \top$ where $v_1 = \text{StrToInt}(s[1 :])$. $\text{ite}(\Phi, v, 0) = v = -v_1$.
- (c) If s contains more than digits (e.g., “abc” and “-abc”), MySQL evaluates non-numerical strings to 0 by default. By the semantics of z3.StrToInt , Φ never holds which leads $\text{ite}(\Phi, v, 0) = 0$.

1671 By 6a, 6c and 6c, we known $\text{ite}(\Phi, s, 0)$ precisely captures the semantics of *SQL's type*

1672 *conversion from strings to integers*.

1673 On the other hand, let us discuss the rule in three cases:

1674 (a) If $\text{StrToInt}(s) \geq 0$, then $v' = \text{StrToInt}(s)$.
 1675 (b) If $\text{StrToInt}(s) < 0$, then $v' = -\text{StrToInt}(s[1:])$.
 1676 (c) If s contains more than digits (e.g., “abc” and “-abc”), MySQL evaluates non-
 1677 numerical strings to 0 by default. By the semantics of this rule, $v' = 0$.
 1678

1679 By 6a, 6c and 6c, we known v' precisely captures the semantics of SQL’s *type conversion*
 1680 *from strings to integers*.

1681 Therefore, $\mathcal{I}(\text{ite}(\Phi, s, 0)) = v'$ and
 1682

$$\begin{aligned} \mathcal{I}(\llbracket \text{StrToInt}(s) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(s = \text{Null}, \text{Null}, \text{ite}(\Phi, v, 0))) \\ &= \text{ite}(s = \text{Null}, \text{Null}, \mathcal{I}(\text{ite}(\Phi, v, 0))) \\ &= \text{ite}(s = \text{Null}, \text{Null}, v') \\ &= \llbracket \text{StrToInt}(s) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1683 7. Inductive case: $E = \text{IntToStr}(v)$
 1684

1685 $\llbracket \text{IntToStr}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null}, \text{Null}, \text{z3.IntToStr}(v))$ by Figure 18.
 1686 $\llbracket \text{IntToStr}(v) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v = \text{Null}, \text{Null}, \text{IntToStr}(v))$ by Figure 17. Note that
 1687 since the Z3 builtin function z3.IntToStr precisely capture the semantics of IntToStr,
 1688 $\mathcal{I}(\text{z3.IntToStr}(v)) = \text{IntToStr}(v)$. Therefore,
 1689

$$\begin{aligned} \mathcal{I}(\llbracket \text{IntToStr}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null}, \text{Null}, \text{z3.IntToStr}(v))) \\ &= \text{ite}(v = \text{Null}, \text{Null}, \mathcal{I}(\text{z3.IntToStr}(v))) \\ &= \text{ite}(v = \text{Null}, \text{Null}, \text{IntToStr}(v)) \\ &= \llbracket \text{IntToStr}(v) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1690 8. Inductive case: $E = \text{DateToStr}(vs)$
 1691

1692 $\llbracket \text{DateToStr}(vs) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(vs = \text{Null}, \text{Null}, y + “-” + m + “-” + d)$ where $y =$
 1693 $\text{z3.IntToStr}(vs[0])$, $m = \text{ite}(vs[1] \leq 9, “0” + \text{z3.IntToStr}(vs[1]), \text{z3.IntToStr}(vs[1]))$,
 1694 and $d = \text{ite}(vs[2] \leq 9, “0” + \text{z3.IntToStr}(vs[2]), \text{z3.IntToStr}(vs[2]))$ by Figure 18.
 1695 $\llbracket \text{DateToStr}(vs) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(vs = \text{Null}, \text{Null}, y' + “-” + m' + “-” + d')$ where
 1696 $y' = \text{IntToStr}(vs[0])$, $m' = \text{ite}(vs[1] \leq 9, “0” + \text{IntToStr}(vs[1]), \text{IntToStr}(vs[1]))$, and
 1697 $d' = \text{ite}(vs[2] \leq 9, “0” + \text{IntToStr}(vs[2]), \text{IntToStr}(vs[2]))$ by Figure 17. Note that since
 1698 the Z3 builtin function z3.IntToStr precisely capture the semantics of IntToStr, $\mathcal{I}(y) = y'$,
 1699 $\mathcal{I}(m) = m'$, and $\mathcal{I}(d) = d'$. Therefore,
 1700

$$\begin{aligned} \mathcal{I}(\llbracket \text{DateToStr}(vs) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(vs = \text{Null}, \text{Null}, y + “-” + m + “-” + d)) \\ &= \text{ite}(vs = \text{Null}, \text{Null}, \mathcal{I}(y) + “-” + \mathcal{I}(m) + “-” + \mathcal{I}(d)) \\ &= \text{ite}(vs = \text{Null}, \text{Null}, y' + “-” + m' + “-” + d') \\ &= \llbracket \text{DateToStr}(v) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1701 9. Inductive case: $E = \text{IntToDate}(v)$
 1702

1703 $\llbracket \text{IntToDate}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null} \vee \neg(\Phi_1 \wedge \Phi_2 \wedge \Phi_3), \text{Null}, [y, m, d])$ where $\text{fdi}(x, y) =$
 1704 $\text{ite}(x \% y = 0, x / y, (x - x \% y) / y)$, $y = \text{fdi}(v, 10^4)$, $m = \text{fdi}(v \% 10^4, 10^2)$, $d = v \% 10^2$,
 1705 $\Phi_0 = y \% 4 = 0 \wedge (y \% 100 \neq 0 \vee y \% 400 = 0)$, $\Phi_1 = \text{MIN_YEAR} \leq y \leq \text{MAX_YEAR}$,
 1706 $\Phi_2 = 1 \leq m \leq 12$, $\Phi_3 = 1 \leq d \wedge (\vee_{c \in \{1, 3, 5, 7, 8, 10, 12\}} m = c \rightarrow d \leq 31) \wedge (m =$
 1707 $2 \rightarrow d \leq 28 + \text{ite}(\Phi_0, 1, 0)) \wedge (\vee_{c \in \{4, 6, 9, 11\}} m = c \rightarrow d \leq 30)$ by Figure 18.
 1708 $\llbracket \text{IntToDate}(v) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null} \vee \text{IsValidDate}(v), \text{Null}, [v'_1, v'_2, v'_3])$ where
 1709 $v'_1 = \lfloor v / 10^4 \rfloor$, $v'_2 = \lfloor (v \% 10^4) / 10^2 \rfloor$, $v'_3 = v \% 10^2$ by Figure 17. By semantics of fdi, we
 1710 know $y = v'_1$, $m = v'_2$ and $d = v'_3$. Note that the function IsValidDate precisely capture the
 1711 semantics of $\neg(\Phi_1 \wedge \Phi_2 \wedge \Phi_3)$, checking whether a date is valid in MySQL. Therefore,
 1712 $\mathcal{I}(\neg(\Phi_1 \wedge \Phi_2 \wedge \Phi_3)) = \text{IsValidDate}(v')$ and
 1713

$$\begin{aligned} \mathcal{I}(\llbracket \text{IntToDate}(v) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null} \vee \neg(\Phi_1 \wedge \Phi_2 \wedge \Phi_3), \text{Null}, [y, m, d])) \\ &= \text{ite}(v = \text{Null} \vee \mathcal{I}(\neg(\Phi_1 \wedge \Phi_2 \wedge \Phi_3)), \text{Null}, \mathcal{I}([y, m, d])) \\ &= \text{ite}(v = \text{Null} \vee \text{IsValidDate}(v'), \text{Null}, [v'_1, v'_2, v'_3]) \\ &= \llbracket \text{IntToDate}(v) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1728 10. Inductive case: $E = \text{StrToDate}(s)$
1729
1730 $\llbracket \text{StrToDate}(s) \rrbracket_{S, \Gamma, \mathcal{T}} = \text{ite}(s = \text{Null}, \text{Null}, \llbracket \text{IntToDate}(v) \rrbracket_{S, \Gamma, \mathcal{T}})$ where
1731 $v = \llbracket \text{StrToInt}(s) \rrbracket_{S, \Gamma, \mathcal{T}}$ by Figure 18. $\llbracket \text{StrToDate}(s) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(s =$
1732 $\text{Null}, \text{Null}, \llbracket \text{IntToDate}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})})$ where $v' = \llbracket \text{StrToInt}(s) \rrbracket_{S, \Gamma, \mathcal{T}}$ by Figure 17. By
1733 inductive hypothesis, we have $\mathcal{I}(\llbracket \text{IntToDate}(v) \rrbracket_{S, \Gamma, \mathcal{T}}) = \llbracket \mathcal{I}(\text{IntToDate}(v)) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} =$
1734 $\llbracket \text{IntToDate}(\mathcal{I}(v)) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \llbracket \text{IntToDate}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$. Therefore,
1735
1736
$$\begin{aligned} \mathcal{I}(\llbracket \text{StrToDate}(s) \rrbracket_{S, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(s = \text{Null}, \text{Null}, \llbracket \text{IntToDate}(v) \rrbracket_{S, \Gamma, \mathcal{T}})) \\ &= \text{ite}(s = \text{Null}, \text{Null}, \mathcal{I}(\llbracket \text{IntToDate}(v) \rrbracket_{S, \Gamma, \mathcal{T}})) \\ &= \text{ite}(s = \text{Null}, \text{Null}, \llbracket \text{IntToDate}(v') \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) \\ &= \llbracket \text{StrToDate}(s) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1737
1738
1739

1740 11. Inductive case: $E = E_1 \diamond E_2$.

1741 Since our extended grammar considers Null, integers, dates and strings, as shown in
1742 Figure 2, the proof for this inductive case is overloaded.

1743
1744 $\llbracket E_1 \diamond E_2 \rrbracket_{S, \Gamma, \mathcal{T}} = \text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, v_1 \diamond v_2)$ where $v_1 = \llbracket \text{ToInt}(E_1) \rrbracket_{S, \Gamma, \mathcal{T}}$
1745 and $v_2 = \llbracket \text{ToInt}(E_2) \rrbracket_{S, \Gamma, \mathcal{T}}$ by Figure 18. $\llbracket E_1 \diamond E_2 \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v'_1 = \text{Null} \vee v'_2 =$
1746 $\text{Null}, \text{Null}, v'_1 \diamond v'_2)$ where $v'_1 = \llbracket \text{ToInt}(E_1) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ and $v'_2 = \llbracket \text{ToInt}(E_2) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$
1747 by Figure 17. By inductive hypothesis, we have $\mathcal{I}(v_1) = \mathcal{I}(\llbracket \text{ToInt}(E_1) \rrbracket_{S, \Gamma, \mathcal{T}}) =$
1748 $\llbracket \text{ToInt}(E_1) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'_1$ and $\mathcal{I}(v_2) = \mathcal{I}(\llbracket \text{ToInt}(E_2) \rrbracket_{S, \Gamma, \mathcal{T}}) = \llbracket \text{ToInt}(E_2) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} =$
1749 v'_2 . Therefore,
1750
1751

$$\begin{aligned} \mathcal{I}(\llbracket E_1 \diamond E_2 \rrbracket_{S, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \text{Null}, v_1 \diamond v_2)) \\ &= \text{ite}(\mathcal{I}(v_1) = \text{Null} \vee \mathcal{I}(v_2) = \text{Null}, \text{Null}, \mathcal{I}(v_1) \diamond \mathcal{I}(v_2)) \\ &= \text{ite}(v'_1 = \text{Null} \vee v'_2 = \text{Null}, \text{Null}, v'_1 \diamond v'_2) \\ &= \llbracket E_1 \diamond E_2 \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1755 12. Inductive case: $E = \text{SubStr}(E_1, E_2, E_3)$.

1756
1757 $\llbracket \text{SubStr}(E_1, E_2, E_3) \rrbracket_{S, \Gamma, \mathcal{T}} = \text{ite}(e_1 = \text{Null} \vee \text{IsStr}(e_2) \vee \text{IsStr}(e_3), \text{Null}, s)$ where
1758 $e_i = \llbracket E_i \rrbracket_{S, \Gamma, \mathcal{T}}$ for $1 \leq i \leq 3$, $e'_1 = \llbracket \text{ToStr}(e_1) \rrbracket_{S, \Gamma, \mathcal{T}}$, $l = \text{z3.Length}(e'_1)$, $e'_2 =$
1759 $\llbracket \text{ToInt}(e_2) \rrbracket_{S, \Gamma, \mathcal{T}}$, $e'_3 = \llbracket \text{ToInt}(e_3) \rrbracket_{S, \Gamma, \mathcal{T}}$, $v = \text{ite}(-l \leq e_2 < 0, \text{ite}(0 < e'_2 \leq l, e'_2 - 1, l +$
1760 $1)), s = \text{ite}(v = 0 \vee v < -l \vee v > l \vee e'_3 \leq 0, \varepsilon, \text{ite}(e'_3 \geq l - v, e'_1[v : l], e'_1[v : 2v + e'_3]))$
1761 by Figure 18.

1762 $\llbracket \text{SubStr}(E_1, E_2, E_3) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(e_4 = \text{Null} \vee \text{IsStr}(e_5) \vee \text{IsStr}(e_6), \text{Null}, s)$ where
1763 $e_{i+3} = \llbracket E_i \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ for $1 \leq i \leq 3$, $e'_4 = \llbracket \text{ToStr}(e_4) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$, $l' = \text{z3.Length}(e'_4)$,
1764 $e'_5 = \llbracket \text{ToInt}(e_5) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$, $e'_6 = \llbracket \text{ToInt}(e_6) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$, $v' = \text{ite}(-l \leq e_5 < 0, \text{ite}(0 <$
1765 $e'_5 \leq l, e'_5 - 1, l + 1)), s' = \text{ite}(v = 0 \vee v < -l \vee v > l \vee e'_6 \leq 0, \varepsilon, \text{ite}(e'_6 \geq l - v, e'_4[v : l], e'_1[v : 2v + e'_6]))$ by Figure 17.

1766 By inductive hypothesis, we have $\mathcal{I}(e_i) = \mathcal{I}(\llbracket E_i \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = \llbracket E_i \rrbracket_{S, \Gamma, \mathcal{T}} = e_{i+3}$ for
1767 $1 \leq i \leq 3$. Then, $\mathcal{I}(e'_1) = \mathcal{I}(\llbracket \text{ToStr}(e_1) \rrbracket_{S, \Gamma, \mathcal{T}}) = \llbracket \text{ToStr}(e_4) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = e'_4$, $\mathcal{I}(e'_2) =$
1768 $\mathcal{I}(\llbracket \text{ToInt}(e_2) \rrbracket_{S, \Gamma, \mathcal{T}}) = \llbracket \text{ToInt}(e_5) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = e'_5$, and $\mathcal{I}(e'_3) = \mathcal{I}(\llbracket \text{ToInt}(e_3) \rrbracket_{S, \Gamma, \mathcal{T}}) =$
1769 $\llbracket \text{ToInt}(e_6) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = e'_6$, $\mathcal{I}(v) = v'$, and $\mathcal{I}(s) = s'$. Furthermore, since the Z3
1770 builtin function z3.Length precisely captures the semantics of len , we have $\mathcal{I}(l) =$
1771 $\mathcal{I}(\text{z3.Length}(e'_1)) = \text{len}(e'_4) = l'$. Therefore,

$$\begin{aligned} \mathcal{I}(\llbracket \text{SubStr}(E_1, E_2, E_3) \rrbracket_{S, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(e_1 = \text{Null} \vee \text{IsStr}(e_2) \vee \text{IsStr}(e_3), \text{Null}, s)) \\ &= \text{ite}(\mathcal{I}(e_1) = \text{Null} \vee \mathcal{I}(\text{IsStr}(e_2)) \vee \mathcal{I}(\text{IsStr}(e_3)), \\ &\quad \text{Null}, \mathcal{I}(s)) \\ &= \text{ite}(e_4 = \text{Null} \vee \text{IsStr}(e_5) \vee \text{IsStr}(e_6), s') \\ &= \llbracket \text{SubStr}(E_1, E_2, E_3) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1778 13. Inductive case: $\phi = \text{Concat}(E_1, E_2)$.

1779
1780 $\llbracket \text{Concat}(E_1, E_2) \rrbracket_{S, \Gamma, \mathcal{T}} = \text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \perp, \text{z3.Concat}(v_1, v_2))$
1781 where $v_1 = \llbracket \text{ToStr}(E_1) \rrbracket_{S, \Gamma, \mathcal{T}}$ and $v_2 = \llbracket \text{ToStr}(E_2) \rrbracket_{S, \Gamma, \mathcal{T}}$ by Figure 18.

1782 $\llbracket \text{Concat}(E_1, E_2) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v'_1 = \text{Null} \vee v'_2 = \text{Null}, \perp, \text{z3.Concat}(v'_1, v'_2))$ where
 1783 $v'_1 = \llbracket \text{ToStr}(E_1) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ and $v'_2 = \llbracket \text{ToStr}(E_2) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By inductive
 1784 hypothesis, we have $\mathcal{I}(v_1) = \mathcal{I}(\llbracket E_1 \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) = \llbracket E_1 \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'_1$ and $\mathcal{I}(v_2) =$
 1785 $\mathcal{I}(\llbracket E_2 \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) = \llbracket E_2 \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'_2$. Furthermore, by the semantics of z3.Concat ,
 1786 $\mathcal{I}(\text{z3.Concat}) = \text{Concat}$. Therefore,
 1787

$$\begin{aligned} \mathcal{I}(\llbracket \text{Concat}(E_1, E_2) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \perp, \text{z3.Concat}(v_1, v_2))) \\ &= \text{ite}(\mathcal{I}(v_1) = \text{Null} \vee \mathcal{I}(v_2) = \text{Null}, \perp, \\ &\quad \mathcal{I}(\text{z3.Concat})(\mathcal{I}(v_1), \mathcal{I}(v_2))) \\ &= \text{ite}(v'_1 = \text{Null} \vee v'_2 = \text{Null}, \perp, \text{Concat}(v'_1, v'_2)) \\ &= \llbracket \text{Concat}(E_1, E_2) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1793 14. Inductive case: $E = \text{Strftime}(\kappa, E)$.

1795 $\llbracket \text{Strftime}(\kappa, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null}, \text{Null}, \text{ite}(\kappa = \text{"%Y"}, v[0], \text{ite}(\kappa =$
 1796 $\text{"%M"}, v[1], v[2])))$ where $v = \llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ by Figure 18. $\llbracket \text{Strftime}(\kappa, E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} =$
 1797 $\text{ite}(v = \text{Null}, \text{Null}, \text{ite}(\kappa = \text{"%Y"}, v[0], \text{ite}(\kappa = \text{"%M"}, v[1], v[2])))$ where $v =$
 1798 $\llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By inductive hypothesis, we have $\mathcal{I}(v) = \mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) =$
 1799 $\llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'$. Therefore,

$$\begin{aligned} \mathcal{I}(\llbracket \text{Strftime}(\kappa, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null}, \text{Null}, \text{ite}(\kappa = \text{"%Y"}, v[0], \\ &\quad \text{ite}(\kappa = \text{"%M"}, v[1], v[2])))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null}, \text{Null}, \text{ite}(\kappa = \text{"%Y"}, \mathcal{I}(v)[0], \\ &\quad \text{ite}(\kappa = \text{"%M"}, \mathcal{I}(v)[1], \mathcal{I}(v)[2]))) \\ &= \text{ite}(v' = \text{Null}, \text{Null}, \text{ite}(\kappa = \text{"%Y"}, v'[0], \\ &\quad \text{ite}(\kappa = \text{"%M"}, v'[1], v'[2]))) \\ &= \llbracket \text{Strftime}(\kappa, E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1808 15. Inductive case: $E = \text{JulianDay}(E)$.

1810 $\llbracket \text{JulianDay}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ToJulianDay}(v)$ where $v = \llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ if v is evaluated to be a
 1811 date by Figure 18. Also, $\text{ToJulianDay}(v) = \lfloor 365.25 * (y + 4716) \rfloor + \lfloor 30.6001 * (m +$
 1812 $4716) \rfloor + d + c - 1524.5$ where $y = v[1] \leq 2?v[0] - 1:v[0]$, $m = v[1] \leq 2?v[1] + 12:v[1]$,
 1813 $d = v[2]$, and $c = 2 - \lfloor y/100 \rfloor + \lfloor y/400 \rfloor$. $\llbracket \text{JulianDay}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = a_1 +$
 1814 $a_2 + d' + c' - 1524.5$ where $v' = \llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$, $y' = \text{ite}(v'[1] \leq 2, v'[0] - 1, v'[0])$,
 1815 $m' = \text{ite}(v'[1] \leq 2, v'[1] + 12, v'[1])$, $d' = v'[2]$, $c' = 2 - \text{fdiv}(y', 100) + \text{fdiv}(y', 400)$,
 1816 $a_1 = \text{fdiv}(36525 * (y' + 4716), 10^2)$, and $a_2 = \text{fdiv}(306001 * (m' + 1), 10^4)$ if v_1 is evaluated
 1817 to be a date by Figure 17. By inductive hypothesis, we have $\mathcal{I}(v) = \mathcal{I}(\llbracket E \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) =$
 1818 $\llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'$. Furthermore, by the semantics of fdiv , $\mathcal{I}(\lfloor 365.25 * (y + 4716) \rfloor) = a_1$
 1819 and $\mathcal{I}(\lfloor 30.6001 * (m + 4716) \rfloor) = a_2$. Therefore,

$$\begin{aligned} \mathcal{I}(\llbracket \text{JulianDay}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ToJulianDay}(v)) \\ &= \mathcal{I}(\lfloor 365.25 * (y + 4716) \rfloor + \lfloor 30.6001 * (m + 4716) \rfloor \\ &\quad + d + c - 1524.5) \\ &= a_1 + a_2 + d' + c' - 1524.5 \\ &= \llbracket \text{JulianDay}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1826 16. Inductive case: $E = \text{DateShift}(E, i, \delta)$.

1827 $\llbracket \text{DateShift}(E, i, \delta) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{DateAdd}(v, i, \delta)$ where $v = \llbracket \text{ToDate}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ if v is eval-
 1828 uated to be a date by Figure 18. Also, $\text{DateAdd}(v, i, \delta)$ is defined as follows:

- If $\delta = \text{"Year"}$, then $\text{DateAdd}(v, i, \delta) = \text{ite}(v'[0] < \text{MIN_YEAR} \vee \text{MAX_YEAR} < v'[0], \text{Null}, \text{Null}, v')$ where $v' = [v[0] + i, v[1], v[2]]$ as i can be negative and dates falling outside the valid date range are regarded as Null.
- If $\delta = \text{"Month"}$, then $\text{DateAdd}(v, i, \delta) = \text{ite}(v'[0] < \text{MIN_YEAR} \vee \text{MAX_YEAR} < v'[0], \text{Null}, \text{Null}, v')$ where $v' = [v[0] + \text{fdiv}(v[1] + i, 12), (v[1] + i) \% 12, v[2]]$.
- If $\delta = \text{"Day"}$, then $\text{DateAdd}(v, i, \delta) = \text{ite}(v' < \text{MIN_DATE} \wedge v' > \text{MAX_DATE}, \text{Null}, v')$ where v' is a new data variable s.t. $\text{SinceBegin}(v') -$

1836 $\text{SinceBegin}(v) = i$. In addition, the SinceBegin function counts the ordinal number
 1837 of a date from a certain day, e.g., “0000-01-01”, which can be defined as
 1838 $\text{SinceBegin}(y, m, d) = \text{year2day}(y) + \text{month2day}(m) + d$ where

$$\begin{aligned} \text{year2day}(y) &= 365 \times y - 1 + \text{fdiv}(y - 1, 4) - \text{fdiv}(y - 1, 100) + \text{fdiv}(y - 1, 400) \\ \text{month2day}(m) &= \sum_{i=1}^{m-1} \text{ite}(i \in \{1, 3, 5, 7, 8, 10, 12\}, 31, \\ &\quad \text{ite}(i = 2, 28 + \text{ite}(\text{leap}(y), 1, 0), 30)) \end{aligned}$$

1843 . Thus, $\text{SinceBegin}(v') - \text{SinceBegin}(v) = i$ ensure the date v' is i days away from
 1844 the date v .

1845 $\llbracket \text{DateShift}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(\delta = \text{“Year”}, \text{DateShiftByYears}(v, i), \text{ite}(\delta = \text{“Month”}, \text{DateShiftByMonths}(v, i), \text{DateShiftByDays}(v, i)))$ where $v_1 = \llbracket E \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ if
 1846 v_1 is evaluated to be a date by Figure 17. By the semantics of the DateShiftByYears ,
 1847 DateShiftByMonths and DateShiftByDays functions, they corresponding to the case 16a,
 1848 16b and 16c. Therefore,

$$\begin{aligned} \mathcal{I}(\llbracket \text{DateShift}(E, i, \delta) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(\delta = \text{“Year”}, \text{DateShiftByYears}(v, i), \\ &\quad \text{ite}(\delta = \text{“Month”}, \text{DateShiftByMonths}(v, i), \\ &\quad \text{DateShiftByDays}(v, i)))) \\ &= \llbracket \text{DateShift}(E, i, \delta) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

□

1855
 1856
 1857
 1858 **Theorem 3** (Correctness of predicate encoding). *Let D be a database over schema \mathcal{S} , xs be a tuple
 1859 list, and ϕ be a predicate. Consider a symbolic database Γ over \mathcal{S} , a list of symbolic tuples \mathcal{T} , and
 1860 ϕ ’s symbolic encoding $\llbracket \phi \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$. For any satisfying interpretation \mathcal{I} with $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs$,
 1861 evaluating ϕ over the database D and the tuple list xs yields the interpretation of ϕ ’s symbolic
 1862 encoding $\mathcal{I}(\llbracket \phi \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$, i.e.,*

$$\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs \Rightarrow \llbracket \phi \rrbracket_{D, xs} = \mathcal{I}(\llbracket \phi \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$$

1863 **Lemma 2.** *Suppose $\llbracket \phi \rrbracket_{D, xs}$ is valid, then $\mathcal{I}(\Gamma) = D \wedge \mathcal{I}(\mathcal{T}) = xs \Rightarrow \llbracket \phi \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \mathcal{I}(\llbracket \phi \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})$ holds.*

1867 *Proof.* Theorem 3 is proved by proving Lemma 2. By structural induction on ϕ .

1868 1. Base cases and some inductive cases are proved in He et al. (2024).

1869 2. Inductive case: $\phi = \text{PrefixOf}(s, E)$.

1870 $\llbracket \text{PrefixOf}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null}, \text{Null}, \text{z3}.\text{PrefixOf}(s, v))$ where
 1871 $v = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ by Figure 19. $\llbracket \text{PrefixOf}(s, E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null}, \text{Null}, \text{PrefixOf}(s, v'))$ where $v' = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By in-
 1872 ductive hypothesis, we have $\mathcal{I}(v) = \mathcal{I}(\llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = v'$.
 1873 Furthermore, since the Z3 builtin function $\text{z3}.\text{PrefixOf}$ precisely captures the semantics
 1874 of PrefixOf , we have $\mathcal{I}(\text{z3}.\text{PrefixOf}(s, v)) = \mathcal{I}(\text{z3}.\text{PrefixOf}(s, \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})) =$
 1875 $\text{PrefixOf}(s, \mathcal{I}(\llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}})) = \text{PrefixOf}(s, \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = \text{PrefixOf}(s, v')$.
 1876 Therefore,

$$\begin{aligned} \mathcal{I}(\llbracket \text{PrefixOf}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null}, \text{Null}, \text{z3}.\text{PrefixOf}(s, v))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null}, \text{Null}, \mathcal{I}(\text{z3}.\text{PrefixOf}(s, v))) \\ &= \text{ite}(v' = \text{Null}, \text{Null}, \text{PrefixOf}(s, v')) \\ &= \llbracket \text{PrefixOf}(s, E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \end{aligned}$$

1880 3. Inductive case: $\phi = \text{SuffixOf}(s, E)$.

1881 $\llbracket \text{SuffixOf}(s, E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null}, \text{Null}, \text{z3}.\text{SuffixOf}(s, v))$ where
 1882 $v = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}}$ by Figure 19. $\llbracket \text{SuffixOf}(s, E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null}, \text{Null}, \text{SuffixOf}(s, v'))$ where $v' = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By in-
 1883 ductive hypothesis, we have $\mathcal{I}(v) = \mathcal{I}(\llbracket \text{ToStr}(E) \rrbracket_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = \llbracket \text{ToStr}(E) \rrbracket_{\mathcal{S}, \Gamma, \mathcal{T}} = v'$.

1890 Furthermore, since the Z3 builtin function `z3.SuffixOf` precisely captures the semantics
 1891 of `SuffixOf`, we have $\mathcal{I}(\text{z3.SuffixOf}(s, v)) = \mathcal{I}(\text{z3.SuffixOf}(s, [\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}})) =$
 1892 $\text{SuffixOf}(s, \mathcal{I}([\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}})) = \text{SuffixOf}(s, [\text{ToStr}(E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = \text{SuffixOf}(s, v')$.
 1893 Therefore,

$$\begin{aligned}\mathcal{I}([\text{SuffixOf}(s, E)]_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null}, \perp, \text{z3.SuffixOf}(s, v))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null}, \perp, \mathcal{I}(\text{z3.SuffixOf}(s, v))) \\ &= \text{ite}(v' = \text{Null}, \perp, \text{SuffixOf}(s, v')) \\ &= [\text{SuffixOf}(s, E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}\end{aligned}$$

1894 4. Inductive case: $\phi = \text{Like}(s, E)$.

1895 $[\text{Like}(s, E)]_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null}, \perp, \text{z3.RegexMatch}(s))$ where $v = [\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}}$
 1896 by Figure 19. $[\text{Like}(s, E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' = \text{Null}, \perp, \text{RegexMatch}(s, v'))$
 1897 where $v' = [\text{ToStr}(E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. By inductive hypothesis, we
 1898 have $\mathcal{I}(v) = \mathcal{I}([\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}}) = [\text{ToStr}(E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = v'$. Furthermore,
 1899 since Z3 precisely support regular expressions, we have $\mathcal{I}(\text{z3.RegexMatch}(s, v)) =$
 1900 $\mathcal{I}(\text{z3.RegexMatch}(s, [\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}})) = \text{RegexMatch}(s, \mathcal{I}([\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}})) =$
 1901 $\text{RegexMatch}(s, [\text{ToStr}(E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = \text{RegexMatch}(s, v')$. Therefore,

$$\begin{aligned}\mathcal{I}([\text{Like}(s, E)]_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null}, \perp, \text{z3.RegexMatch}(s, v))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null}, \perp, \mathcal{I}(\text{z3.RegexMatch}(s, v))) \\ &= \text{ite}(v' = \text{Null}, \perp, \text{RegexMatch}(s, v')) \\ &= [\text{Like}(s, E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}\end{aligned}$$

1902 5. Inductive case: $\phi = \text{Contain}(s, E)$.

1903 $[\text{Contain}(s, E)]_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v = \text{Null}, \perp, \text{z3.RegexMatch}(s', v))$ where
 1904 $s' = \text{Concat}(\text{“*”}, s, \text{“*”})$ and $v = [\text{ToStr}(E)]_{\mathcal{S}, \Gamma, \mathcal{T}}$ by Figure 19.
 1905 $[\text{Contain}(s, E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = [\text{Like}(s'', E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v' =$
 1906 $\text{Null}, \perp, \text{RegexMatch}(s'', v'))$ where $s'' = \text{Concat}(\text{“%”}, s, \text{“%”})$ and $v' =$
 1907 $[\text{ToStr}(E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ by Figure 17. Furthermore, by the semantics of `z3.RegexMatch`
 1908 and `RegexMatch`, we know s' and s'' represent the same regular expression, and
 1909 $\mathcal{I}(\text{z3.RegexMatch}(s', x')) = \text{RegexMatch}(s'', x'')$ iff $\mathcal{I}(x') = x''$. Therefore,

$$\begin{aligned}\mathcal{I}([\text{Contain}(s, E)]_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v = \text{Null}, \perp, \text{z3.RegexMatch}(s', v))) \\ &= \text{ite}(\mathcal{I}(v) = \text{Null}, \perp, \mathcal{I}(\text{z3.RegexMatch}(s', v))) \\ &= \text{ite}(v' = \text{Null}, \perp, \text{RegexMatch}(s'', v')) \\ &= [\text{Like}(s'', E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} \\ &= [\text{Contain}(s, E)]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}\end{aligned}$$

1910 6. Inductive case: $\phi = E_1 \odot E_2$.

1911 $[E_1 \odot E_2]_{\mathcal{S}, \Gamma, \mathcal{T}} = \text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \perp, v_1 \odot v_2)$ where $v_1 = [E_1]_{\mathcal{S}, \Gamma, \mathcal{T}}$ and
 1912 $v_2 = [E_2]_{\mathcal{S}, \Gamma, \mathcal{T}}$ if v_1 and v_2 share the same type, i.e., $\text{Type}(v_1) = \text{Type}(v_2)$ by Figure 19.
 1913 $[E_1 \odot E_2]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})} = \text{ite}(v'_1 = \text{Null} \vee v'_2 = \text{Null}, \perp, v'_1 \odot v'_2)$ where $v_1 = [E_1]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$
 1914 and $v'_2 = [E_2]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}$ if v'_1 and v'_2 share the same type, i.e., $\text{Type}(v'_1) = \text{Type}(v'_2)$ by
 1915 Figure 17. Note that this operation only works for E_1 and E_2 sharing the same type which
 1916 is consistent with MySQL. By inductive hypothesis, we have $\mathcal{I}(v_1) = \mathcal{I}([E_1]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) =$
 1917 $[E_1]_{\mathcal{S}, \Gamma, \mathcal{T}} = v'_1$ and $\mathcal{I}(v_2) = \mathcal{I}([E_2]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}) = [E_2]_{\mathcal{S}, \Gamma, \mathcal{T}} = v'_2$. Therefore, when E_1
 1918 and E_2 have the same type, we have

$$\begin{aligned}\mathcal{I}([E_1 \odot E_2]_{\mathcal{S}, \Gamma, \mathcal{T}}) &= \mathcal{I}(\text{ite}(v_1 = \text{Null} \vee v_2 = \text{Null}, \perp, v_1 \odot v_2)) \\ &= \text{ite}(\mathcal{I}(v_1) = \text{Null} \vee \mathcal{I}(v_2) = \text{Null}, \perp, \mathcal{I}(v_1) \odot \mathcal{I}(v_2)) \\ &= \text{ite}(v'_1 = \text{Null} \vee v'_2 = \text{Null}, \perp, v'_1 \odot v'_2) \\ &= [E_1 \odot E_2]_{\mathcal{I}(\Gamma), \mathcal{I}(\mathcal{T})}\end{aligned}$$

1919 \square

1920 **Theorem 2** (Equivalence under set semantics). *Given two relations $R_1 = [t_1, \dots, t_n]$ and $R_2 = [r_1, \dots, r_m]$, if formula (2) is valid, then R_1 and R_2 are equivalent under set semantics.*

1944 *Proof.* Let F_1 be the first conjunct of formula (2), i.e., $\bigwedge_{i=1}^n (\neg \text{Del}(t_i) \rightarrow \bigvee_{j=1}^m (\neg \text{Del}(r_j) \wedge t_i = r_j))$,
 1945 and let F_2 be the second conjunct of formula (2), i.e., $\bigwedge_{j=1}^m (\neg \text{Del}(r_j) \rightarrow \bigvee_{i=1}^n (\neg \text{Del}(t_i) \wedge r_j = t_i))$.
 1946 Since formula (2) is valid, both F_1 and F_2 are valid. Now consider F_1 . It specifies for any tuple
 1947 $t_i \in R_1$, if t_i is not deleted, then there exists a tuple r_j that is not deleted and $t_i = r_j$. By the
 1948 definition of \subseteq , $R_1 \subseteq R_2$. Similarly, F_2 specifies $R_2 \subseteq R_1$. Therefore, $R_1 = R_2$.
 1949

□

1950
 1951
 1952
 1953
 1954
 1955
 1956
 1957
 1958
 1959
 1960
 1961
 1962
 1963
 1964
 1965
 1966
 1967
 1968
 1969
 1970
 1971
 1972
 1973
 1974
 1975
 1976
 1977
 1978
 1979
 1980
 1981
 1982
 1983
 1984
 1985
 1986
 1987
 1988
 1989
 1990
 1991
 1992
 1993
 1994
 1995
 1996
 1997