

RESDSQL: Decoupling Schema Linking and Skeleton Parsing for Text-to-SQL

Haoyang Li^{1,2,3}, Jing Zhang^{1,2,3*}, Cuiping Li^{1,2,3}, Hong Chen^{1,2,3}

¹ Key Laboratory of Data Engineering and Knowledge Engineering of Ministry of Education, Renmin University of China

² Engineering Research Center of Ministry of Education on Database and BI

³ Information School, Renmin University of China

{lihaoyang.cs, zhang-jing, licuiping, chong}@ruc.edu.cn

Abstract

One of the recent best attempts at Text-to-SQL is the pre-trained language model. Due to the structural property of the SQL queries, the seq2seq model takes the responsibility of parsing both the schema items (*i.e.*, tables and columns) and the skeleton (*i.e.*, SQL keywords). Such coupled targets increase the difficulty of parsing the correct SQL queries especially when they involve many schema items and logic operators. This paper proposes a ranking-enhanced encoding and skeleton-aware decoding framework to decouple the schema linking and the skeleton parsing. Specifically, for a seq2seq encoder-decode model, its encoder is injected by the most relevant schema items instead of the whole unordered ones, which could alleviate the schema linking effort during SQL parsing, and its decoder first generates the skeleton and then the actual SQL query, which could implicitly constrain the SQL parsing. We evaluate our proposed framework on Spider and its three robustness variants: Spider-DK, Spider-Syn, and Spider-Realistic. The experimental results show that our framework delivers promising performance and robustness. Our code is available at <https://github.com/RUCKBReasoning/RESDSQL>.

Introduction

Relational databases that are used to store heterogeneous data types including text, integer, float, etc., are omnipresent in modern data management systems. However, ordinary users usually cannot make the best use of databases because they are not good at translating their requirements to the database language—*i.e.*, the structured query language (SQL). To assist these non-professional users in querying the databases, researchers propose the Text-to-SQL task (Yu et al. 2018a; Cai et al. 2018), which aims to automatically translate users’ natural language questions into SQL queries. At the same time, related benchmarks are becoming increasingly complex, from the single-domain benchmarks such as ATIS (Iyer et al. 2017) and GeoQuery (Zelle and Mooney 1996) to the cross-domain benchmarks such as WikiSQL (Zhong, Xiong, and Socher 2017) and Spider (Yu et al. 2018c). Most of the recent works are done on Spider because it is the most challenging benchmark which in-

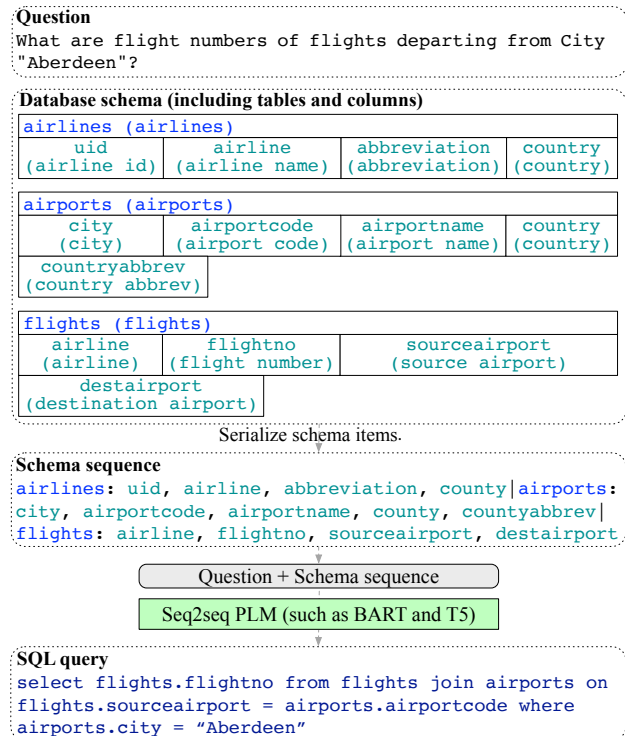


Figure 1: Illustration of a Text-to-SQL instance solved by a seq2seq PLM. In the database schema, each schema item is denoted by its “original name (semantic name)”.

volves many complex SQL operators (such as GROUP BY, ORDER BY, and HAVING, etc.) and nested SQL queries.

With the recent advances in pre-trained language models (PLMs), many existing works formulate the Text-to-SQL task as a semantic parsing problem and use a sequence-to-sequence (seq2seq) model to solve it (Scholak, Schucher, and Bahdanau 2021; Shi et al. 2021; Shaw et al. 2021). Concretely, as shown in Figure 1, given a question and a database schema, the schema items are serialized into a schema sequence where the order of the schema items is either default or random. Then, a seq2seq PLM, such as BART (Lewis et al. 2020) and T5 (Raffel et al. 2020), is leveraged to generate the SQL query based on the concatenation of the ques-

*Jing Zhang is the corresponding author.

tion and the schema sequence. We observe that the target SQL query contains not only the skeleton that reveals the logic of the question but also the required schema items. For instance, for a SQL query: “SELECT petid FROM pets WHERE pet_age = 1”, its skeleton is “SELECT _ FROM _ WHERE _” and its required schema items are “petid”, “pets”, and “pet_age”.

Since Text-to-SQL needs to perform not only the schema linking which aligns the mentioned entities in the question to schema items in the database schema, but also the skeleton parsing which parses out the skeleton of the SQL query, the major challenges are caused by a large number of required schema items and the complex composition of operators such as GROUP BY, HAVING, and JOIN ON involved in a SQL query. The intertwining of the schema linking and the skeleton parsing complicates learning even more.

To investigate whether the Text-to-SQL task could become easier if the schema linking and the skeleton parsing are decoupled, we conduct a preliminary experiment on Spider’s dev set. Concretely, we fine-tune a T5-Base model to generate the pure skeletons based on the questions (*i.e.*, skeleton parsing task). We observe that the exact match accuracy on such a task achieves about 80% using the fine-tuned T5-Base. However, even the T5-3B model only achieves about 70% accuracy (Shaw et al. 2021; Scholak, Schucher, and Bahdanau 2021). This pre-experiment indicates that decoupling such two objectives could be a potential way of reducing the difficulty of Text-to-SQL.

To realize the above decoupling idea, we propose a Ranking-enhanced Encoding plus a Skeleton-aware Decoding framework for Text-to-SQL (RESDSL). The former injects a few but most relevant schema items into the seq2seq model’s encoder instead of all schema items. In other words, the schema linking is conducted beforehand to filter out most of the irrelevant schema items in the database schema, which can alleviate the difficulty of the schema linking for the seq2seq model. For such purpose, we train an additional cross-encoder to classify the tables and columns simultaneously based on the input question, and then rank and filter them according to the classification probabilities to form a ranked schema sequence. The latter does not add any new modules but simply allows the seq2seq model’s decoder to first generate the SQL skeleton, and then the actual SQL query. Since skeleton parsing is much easier than SQL parsing, the first generated skeleton could implicitly guide the subsequent SQL parsing via the masked self-attention mechanism in the decoder.

Contributions (1) We investigate a potential way of decoupling the schema linking and the skeleton parsing to reduce the difficulty of Text-to-SQL. Specifically, we propose a ranking-enhanced encoder to alleviate the effort of the schema linking and a skeleton-aware decoder to implicitly guide the SQL parsing by the skeleton. (2) We conduct extensive evaluation and analysis and show that our framework not only achieves the new state-of-the-art (SOTA) performance on Spider but also exhibits strong robustness.

Problem Definition

Database Schema A relational database is denoted as \mathcal{D} . The database schema \mathcal{S} of \mathcal{D} includes (1) a set of N tables $\mathcal{T} = \{t_1, t_2, \dots, t_N\}$, (2) a set of columns $\mathcal{C} = \{c_1^1, \dots, c_{n_1}^1, c_1^2, \dots, c_{n_2}^2, \dots, c_1^N, \dots, c_{n_N}^N\}$ associated with the tables, where n_i is the number of columns in the i -th table, (3) and a set of foreign key relations $\mathcal{R} = \{(c_k^i, c_h^j) | c_k^i, c_h^j \in \mathcal{C}\}$, where each (c_k^i, c_h^j) denotes a foreign key relation between column c_k^i and column c_h^j . We use $M = \sum_{i=1}^N n_i$ to denote the total number of columns in \mathcal{D} .

Original Name and Semantic Name We use “schema items” to uniformly refer to tables and columns in the database. Each schema item can be represented by an original name and a semantic name. The semantic name can indicate the semantics of the schema item more precisely. As shown in Figure 1, it is obvious that the semantic names “airline id” and “destination airport” are more clear than their original names “uid” and “destairport”. Sometimes the semantic name is the same as the original name.

Text-to-SQL Task Formally, given a question q in natural language and a database \mathcal{D} with its schema \mathcal{S} , the Text-to-SQL task aims to translate q into a SQL query l that can be executed on \mathcal{D} to answer the question q .

Methodology

In this section, we first give an overview of the proposed framework and then delve into its design details.

Model Overview

Following Shaw et al. (2021); Scholak, Schucher, and Bahdanau (2021), we treat Text-to-SQL as a translation task, which can be solved by an encoder-decoder transformer model. Facing the above problems, we extend the existing seq2seq Text-to-SQL methods by injecting the most relevant schema items in the input sequence and the SQL skeleton in the output sequence, which results in a ranking-enhanced encoder and a skeleton-aware decoder. We provide the high-level overview of the proposed RESDSL framework in Figure 2. The encoder of the seq2seq model receives the ranked schema sequence, such that the schema linking effort could be alleviated during SQL parsing. To obtain such a ranked schema sequence, an additional cross-encoder is proposed to classify the schema items according to the given question, and then we rank and filter them based on the classification probabilities. The decoder of the seq2seq model first parses out the SQL skeleton and then the actual SQL query, such that the SQL generation can be implicitly constrained by the previously parsed skeleton. By doing this, to a certain extent, the schema linking and the skeleton parsing are not intertwined but decoupled.

Ranking-Enhanced Encoder

Instead of injecting all schema items, we only consider the most relevant schema items in the input of the encoder. For this purpose, we devise a cross-encoder to classify the tables and columns simultaneously and then rank them based on

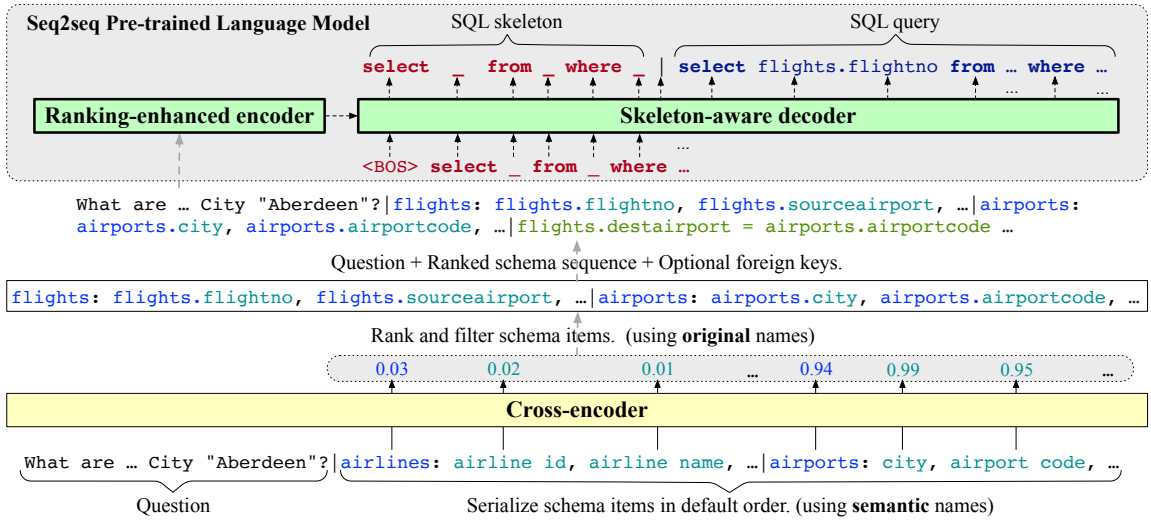


Figure 2: An overview of the ranking-enhanced encoding and skeleton-aware decoding framework. We train a cross-encoder for classifying the schema items. Then we take the question, the ranked schema sequence, and optional foreign keys as the input of the ranking-enhanced encoder. The skeleton-aware decoder first decodes the SQL skeleton and then the actual SQL query.

their probabilities. Based on the ranking order, on one hand, we filter out the irrelevant schema items. On the other hand, we use the ranked schema sequence instead of the unordered schema sequence, so that the seq2seq model could capture potential position information for schema linking.

As for the input of the cross-encoder, we flatten the schema items into a schema sequence in their *default* order and concatenate it with the question to form an input sequence: $X = q | t_1 : c_1^1, \dots, c_{n_1}^1 | \dots | t_N : c_1^N, \dots, c_{n_N}^N$, where $|$ is the delimiter. To better represent the semantics of schema items, instead of their original names, we use their semantic names which are closer to the natural expression.

Encoding Module We feed X into RoBERTa (Liu et al. 2019), an improved version of BERT (Devlin et al. 2019). Since each schema item will be tokenized into one or more tokens by PLM’s tokenizer (e.g., the column “airline id” will be split into two tokens: “airline” and “id”), and our target is to represent each schema item as a whole for classification, we need to pool the output embeddings belonging to each schema item. To achieve this goal, we use a pooling module that consists of a two-layer BiLSTM (Hochreiter and Schmidhuber 1997) and a non-linear fully-connected layer. After pooling, each table embedding can be denoted by $\mathbf{T}_i \in \mathbb{R}^{1 \times d}$ ($i \in \{1, \dots, N\}$) and each column embedding can be denoted by $\mathbf{C}_k^i \in \mathbb{R}^{1 \times d}$ ($i \in \{1, \dots, N\}, k \in \{1, \dots, n_i\}$), where d denotes the hidden size.

Column-Enhanced Layer We observe that some questions only mention the column name rather than the table name. For example in Figure 1, the question mentions a column name “city”, but its corresponding table name “airports” is ignored. This table name missing issue may compromise the table classification performance. Therefore, we propose a column-enhanced layer to inject the column information into the corresponding table embedding. In this way,

a table could be identified even if the question only mentions its columns. Concretely, for the i -th table, we inject the column information $\mathbf{C}_k^i \in \mathbb{R}^{n_i \times d}$ into the table embedding \mathbf{T}_i by stacking a multi-head scaled dot-product attention layer (Vaswani et al. 2017) and a feature fusion layer on the top of the encoding module:

$$\begin{aligned} \mathbf{T}_i^C &= \text{MultiHeadAttn}(\mathbf{T}_i, \mathbf{C}_k^i, \mathbf{C}_k^i, h), \\ \hat{\mathbf{T}}_i &= \text{Norm}(\mathbf{T}_i + \mathbf{T}_i^C). \end{aligned} \quad (1)$$

Here, \mathbf{T}_i acts as the query and \mathbf{C}_k^i acts as both the key and the value, h is the number of heads, and $\text{Norm}(\cdot)$ is a row-wise L_2 normalization function. \mathbf{T}_i^C represents the column-attentive table embedding. We fuse the original table embedding \mathbf{T}_i and the column-attentive table embedding \mathbf{T}_i^C to obtain the column-enhanced table embedding $\hat{\mathbf{T}}_i \in \mathbb{R}^{1 \times d}$.

Loss Function of Cross-Encoder Cross-entropy loss is a well-adopted loss function in classification tasks. However, since a SQL query usually involves only a few tables and columns in the database, the label distribution of the training set is highly imbalanced. As a result, the number of negative examples is many times that of positive examples, which will induce serious training bias. To alleviate this issue, we employ the focal loss (Lin et al. 2017) as our classification loss. Then, we form the loss function of the cross-encoder in a multi-task learning way, which consists of both the table classification loss and the column classification loss, *i.e.*,

$$\mathcal{L}_1 = \frac{1}{N} \sum_{i=1}^N FL(y_i, \hat{y}_i) + \frac{1}{M} \sum_{i=1}^N \sum_{k=1}^{n_i} FL(y_k^i, \hat{y}_k^i), \quad (2)$$

where FL denotes the focal loss function and y_i is the ground truth label of the i -th table. $y_i = 1$ indicates the table is referenced by the SQL query and 0 otherwise. y_k^i is the ground truth label of the k -th column in the i -th table.

Similarly, $y_k^i = 1$ indicates the column is referenced by the SQL query and 0 otherwise. \hat{y}_i and \hat{y}_k^i are predicted probabilities, which are estimated by two different MLP modules based on the table and column embeddings \hat{T}_i and C_k^i :

$$\begin{aligned}\hat{y}_i &= \sigma((\hat{T}_i U_1^t + \mathbf{b}_1^t) U_2^t + \mathbf{b}_2^t), \\ \hat{y}_k^i &= \sigma((C_k^i U_1^c + \mathbf{b}_1^c) U_2^c + \mathbf{b}_2^c),\end{aligned}\quad (3)$$

where $U_1^t, U_1^c \in \mathbb{R}^{d \times w}$, $\mathbf{b}_1^t, \mathbf{b}_1^c \in \mathbb{R}^w$, $U_2^t, U_2^c \in \mathbb{R}^{w \times 2}$, $\mathbf{b}_2^t, \mathbf{b}_2^c \in \mathbb{R}^2$ are trainable parameters, and $\sigma(\cdot)$ denotes Softmax.

Prepare Input for Ranking-Enhanced Encoder During inference, for each Text-to-SQL instance, we leverage the above-trained cross-encoder to compute a probability for each schema item. Then, we only keep top- k_1 tables in the database and top- k_2 columns for each remained table to form a ranked schema sequence. k_1 and k_2 are two important hyper-parameters. When k_1 or k_2 is too small, a portion of the required tables or columns may be excluded, which is fatal for the subsequent seq2seq model. As k_1 or k_2 becomes larger, more and more irrelevant tables or columns may be introduced as noise. Therefore, we need to choose appropriate values for k_1 and k_2 to ensure a high recall while preventing the introduction of too much noise. The input sequence for the ranking-enhanced encoder (*i.e.*, seq2seq model’s encoder) is formed as the concatenation of the question, the ranked schema sequence, and optional foreign key relations (see Figure 2). Foreign key relations contain rich information about the structure of the database, which could promote the generation of the JOIN ON clauses. In the ranked schema sequence, we use the original names instead of the semantic names because the schema items in the SQL queries are represented by their original names, and using the former will facilitate the decoder to *directly copy required schema items* from the input sequence.

Skeleton-Aware Decoder

Most seq2seq Text-to-SQL methods tell the decoder to generate the target SQL query directly. However, the apparent gap between the natural language and the SQL query makes it difficult to perform the correct generation. To alleviate this problem, we would like to decompose the SQL generation into two steps: (1) generate the SQL skeleton based on the semantics of the question, and then (2) select the required “data” (*i.e.*, tables, columns, and values) from the input sequence to fill the slots in the skeleton.

To realize the above decomposition idea without adding additional modules, we propose a new generation objective based on the intrinsic characteristic of the transformer decoder, which generates the t -th token depending on not only the output of the encoder but also the output of the decoder before the t -th time step (Vaswani et al. 2017). Concretely, instead of decoding the target SQL directly, we encourage the decoder to first decode the skeleton of the SQL query, and based on this, we continue to decode the SQL query.

By parsing the skeleton first and then parsing the SQL query, at each decoding step, SQL generation will be easier because the decoder could either copy a “data” from the input sequence or a SQL keyword from the previously parsed

Q	List the duration, file size and format of songs whose genre is pop, ordered by title?
SQL _o	SELECT T1.duration, T1.file_size, T1.formats FROM files AS T1 JOIN song AS T2 ON T1.f_id = T2.f_id WHERE T2.genre_is = “pop” ORDER BY T2.song_name
SQL _n	select files.duration, files.file_size, files.formats from files join song on files.f_id = song.f_id where song.genre_is = ‘pop’ order by song.song_name asc
SQL _s	select _ from _ where _ order by _ asc

Table 1: An example from Spider. Here, Q, SQL_o, SQL_n, and SQL_s denote the question, the original SQL query, the normalized SQL query, and the SQL skeleton, respectively.

skeleton. Now, the objective of the seq2seq model is:

$$\mathcal{L}_2 = \frac{1}{G} \sum_{i=1}^G p(l_i^s, l_i | S_i), \quad (4)$$

where G is the number of Text-to-SQL instances, S_i is the input sequence of the i -th instance which consists of the question, the ranked schema sequence, and optional foreign key relations. l_i denotes the i -th target SQL query and l_i^s is the skeleton extracted from l_i . We will present some necessary details on how to normalize SQL queries and how to extract their skeletons.

SQL Normalization The Spider dataset is manually created by 11 annotators with different annotation habits, which results in slightly different styles among the final annotated SQL queries, such as uppercase versus lowercase keywords. Although different styles have no impact on the execution results, the model requires some extra effort to learn and adapt to them. To reduce the learning difficulty, we normalize the original SQL queries before training by (1) unifying the keywords and schema items into lowercase, (2) adding spaces around parentheses and replacing double quotes with single quotes, (3) adding an ASC keyword after the ORDER BY clause if it does not specify the order, and (4) removing the AS clause and replacing all table aliases with their original names. We present an example in Table 1.

SQL Skeleton Extraction Based on the normalized SQL queries, we can extract their skeletons which only contain SQL keywords and slots. Specifically, given a normalized SQL query, we keep its keywords and replace the rest parts with slots. Note that we do not keep the JOIN ON keyword because it is difficult to find a counterpart from the question (Gan et al. 2021b). As shown in Table 1, although the original SQL query looks complex, its skeleton is simple and each keyword can find a counterpart from the question. For example, “order by _ asc” in the skeleton can be inferred from “ordered by title?” in the question.

Execution-Guided SQL Selector Since we do not constrain the decoder with SQL grammar, the model may generate some illegal SQL queries. To alleviate this problem, we follow Suhr et al. (2020) to use an execution-guided SQL selector which performs the beam search during the decoding

procedure and then selects the first executable SQL query in the beam as the final result.

Experiments

Experimental Setup

Datasets We conduct extensive experiments on Spider and its three variants which are proposed to evaluate the robustness of the Text-to-SQL parser. Spider (Yu et al. 2018c) is the most challenging benchmark for the cross-domain and multi-table Text-to-SQL task. Spider contains a training set with 7,000 samples¹, a dev set with 1,034 samples, and a hidden test set with 2,147 samples. There is no overlap between the databases in different splits. For robustness, we train the model on Spider’s training set but evaluate it on Spider-DK (Gan, Chen, and Purver 2021) with 535 samples, Spider-Syn (Gan et al. 2021a) with 1034 samples, and Spider-Realistic (Deng et al. 2021) with 508 samples. These evaluation sets are derived from Spider by modifying questions to simulate real-world application scenarios. Concretely, Spider-DK incorporates some domain knowledge to paraphrase questions. Spider-Syn replaces schema-related words with synonyms in questions. Spider-Realistic removes explicitly mentioned column names in questions.

Evaluation Metrics To evaluate the performance of the Text-to-SQL parser, following Yu et al. 2018c; Zhong, Yu, and Klein 2020, we adopt two metrics: Exact-set-Match accuracy (EM) and EXecution accuracy (EX). The former measures whether the predicted SQL query can be exactly matched with the gold SQL query by converting them into a special data structure (Yu et al. 2018c). The latter compares the execution results of the predicted SQL query and the gold SQL query. The EX metric is sensitive to the generated values, but the EM metric is not. In practice, we use the sum of EM and EX to select the best checkpoint of the seq2seq model. For the cross-encoder, we use Area Under ROC Curve (AUC) to evaluate its performance. Since the cross-encoder classifies tables and columns simultaneously, we adopt the sum of table AUC and column AUC to select the best checkpoint of the cross-encoder.

Implementation Details We train RESDSQL in two stages. In the first stage, we train the cross-encoder for ranking schema items. The number of heads h in the column-enhanced layer is 8. We use AdamW (Loshchilov and Hutter 2019) with batch size 32 and learning rate $1e-5$ for optimization. In the focal loss, the focusing parameter γ and the weighted factor α are set to 2 and 0.75 respectively. Then, k_1 and k_2 are set to 4 and 5 according to the statistics of the datasets. For training the seq2seq model in the second stage, we consider three scales of T5: Base, Large, and 3B. We fine-tune them with Adafactor (Shazeer and Stern 2018) using different batch size (bs) and learning rate (lr), resulting in RESDSQL-Base (bs = 32, lr = $1e-4$), RESDSQL-Large (bs

¹Spider also provides additional 1,659 training samples, which are collected from some single-domain datasets, such as GeoQuery (Zelle and Mooney 1996) and Restaurants (Giordani and Moschitti 2012). But following (Scholak, Schucher, and Bahdanau 2021), we ignore this part in our training set.

= 32, lr = $5e-5$), and RESDSQL-3B (bs = 96, lr = $5e-5$). For both stages of training, we adopt linear warm-up (the first 10% training steps) and cosine decay to adjust the learning rate. We set the beam size to 8 during decoding. Moreover, following Lin, Socher, and Xiong (2020), we extract potentially useful contents from the database to enrich the column information.

Environments We conduct all experiments on a server with one NVIDIA A100 (80G) GPU, one Intel(R) Xeon(R) Silver 4316 CPU, 256 GB memory and Ubuntu 20.04.2 LTS operating system.

Results on Spider

Table 2 reports EM and EX results on Spider. Noticeably, we observe that RESDSQL-Base achieves better performance than the bare T5-3B, which indicates that our decoupling idea can substantially reduce the learning difficulty of Text-to-SQL. Then, RESDSQL-3B outperforms the best baseline by 1.6% EM and 1.3% EX on the dev set. Furthermore, when combined with NatSQL (Gan et al. 2021b), an intermediate representation of SQL, RESDSQL-Large achieves competitive results compared to powerful baselines on the dev set, and RESDSQL-3B achieves new SOTA performance on both the dev set and the test set. Specifically, on the dev set, RESDSQL-3B + NatSQL brings 4.2% EM and 3.6% EX absolute improvements. On the hidden test set, RESDSQL-3B + NatSQL achieves competitive performance on EM and dramatically increases EX from 75.5% to 79.9% (+4.4%), showing the effectiveness of our approach. The reason for the large gap between EM (72.0%) and EX (79.9%) is that EM is overly strict (Zhong, Yu, and Klein 2020). For example in Spider, given a question “Find id of the candidate who most recently accessed the course?”, its gold SQL query is “select candidate_id from candidate_assessments order by assessment_date desc limit 1”. In fact, there is another SQL query “select candidate_id from candidate_assessments where assessment_date = (select max(assessment_date) from candidate_assessments)” which can also be executed to answer the question (*i.e.*, EX is positive). However, EM will judge the latter to be wrong, which leads to false negatives.

Results on Robustness Settings

Recent studies (Gan et al. 2021a; Deng et al. 2021) show that neural Text-to-SQL parsers are fragile to question perturbations because explicitly mentioned schema items are removed or replaced with semantically consistent words (*e.g.*, synonyms), which increases the difficulty of schema linking. Therefore, more and more efforts have been recently devoted to improving the robustness of neural Text-to-SQL parsers, such as TTK (Qin et al. 2022) and SUN (Gao et al. 2022). To validate the robustness of RESDSQL, we train our model on Spider’s training set and evaluate it on three challenging Spider variants: Spider-DK, Spider-Syn, and Spider-Realistic. Results are reported in Table 3. We can observe that in all three datasets, RESDSQL-3B + NatSQL surprisingly outperforms all strong competitors by a large margin, which suggests that our decoupling idea can also improve

Approach	Dev Set		Test Set	
	EM	EX	EM	EX
Non-seq2seq methods				
RAT-SQL + GRAPPA (Yu et al. 2021)	73.4	-	69.6	-
RAT-SQL + GAP + NatSQL (Gan et al. 2021b)	73.7	75.0	68.7	73.3
SMBOP + GRAPPA (Rubin and Berant 2021)	74.7	75.0	69.5	71.1
DT-Fixup SQL-SP + RoBERTa (Xu et al. 2021)	75.0	-	70.9	-
LGESQL + ELECTRA (Cao et al. 2021)	75.1	-	72.0	-
S ² SQL + ELECTRA (Hui et al. 2022)	76.4	-	72.1	-
Seq2seq methods				
T5-3B (Scholak, Schucher, and Bahdanau 2021)	71.5	74.4	68.0	70.1
T5-3B + PICARD (Scholak, Schucher, and Bahdanau 2021)	75.5	79.3	71.9	75.1
RASAT + PICARD (Qi et al. 2022)	75.3	80.5	70.9	75.5
Our proposed method				
RESDSL-Base	71.7	77.9	-	-
RESDSL-Base + NatSQL	74.1	80.2	-	-
RESDSL-Large	75.8	80.1	-	-
RESDSL-Large + NatSQL	76.7	81.9	-	-
RESDSL-3B	78.0	81.8	-	-
RESDSL-3B + NatSQL	80.5	84.1	72.0	79.9

Table 2: EM and EX results on Spider’s development set and hidden test set (%). We compare our approach with some powerful baseline methods from the top of the official leaderboard of Spider.

the robustness of seq2seq Text-to-SQL parsers. We attribute this to the fact that our proposed cross-encoder can alleviate the difficulty of schema linking and thus exhibits robustness in terms of question perturbations.

Ablation Studies

We take a thorough ablation study on Spider’s dev set to analyze the effect of each design.

Effect of Column-Enhanced Layer We investigate the effectiveness of the column-enhanced layer, which is designed to alleviate the table missing problem. Table 4 shows that removing such a layer will lead to a decrease in the total AUC, as it can inject the human prior into the cross-encoder.

Effect of Focal Loss We also study the effect of focal loss by replacing it with the cross-entropy loss for schema item classification. Table 4 shows that cross-entropy leads to a performance drop because it cannot alleviate the label-imbalance problem in the training data.

Effect of Ranking Schema Items As shown in Table 5, when we replace the ranked schema sequence with the original unordered schema sequence, EM and EX significantly decrease by 4.5% and 7.8% respectively. This result proves that the ranking-enhanced encoder takes a crucial role.

Effect of Skeleton Parsing Meanwhile, from Table 5, we can observe that EM and EX drop 0.7% and 0.8% respectively when removing the SQL skeleton from the decoder’s output (*i.e.*, without skeleton parsing). This is because the seq2seq model needs to make extra efforts to bridge the gap between natural language questions and SQL queries when parsing SQL queries directly.

Related Work

Our method is related to the encoder-decoder architecture designed for Text-to-SQL, the schema item classification task, and the intermediate representation.

Encoder-Decoder Architecture

The encoder aims to jointly encode the question and database schema, which is generally divided into sequence encoder and graph encoder. The decoder aims to generate the SQL queries based on the output of the encoder. Due to the special format of SQL, grammar- and execution-guided decoders are studied to constrain the decoding results.

Sequence Encoder The input is a sequence that concatenates the question with serialized database schema (Yu et al. 2021; Lin, Socher, and Xiong 2020). Then, each token in the sequence is encoded by a PLM encoder, such as BERT (Devlin et al. 2019) and encoder part of T5 (Raffel et al. 2020).

Graph Encoder The input is one or more heterogeneous graphs (Wang et al. 2020a; Hui et al. 2022; Cao et al. 2021; Cai et al. 2021), where a node represents a question token, a table or a column, and an edge represents the relation between two nodes. Then, relation-aware transformer networks (Shaw, Uszkoreit, and Vaswani 2018) or relational graph neural networks, such as RGCN (Schlichtkrull et al. 2018) and RGAT (Wang et al. 2020b), are applied to encode each node. Some works also employ PLM encoders to initialize the representation of nodes on the graph (Cao et al. 2021; Wang et al. 2020a; Rubin and Berant 2021). It is undeniable that the graph encoder can flexibly and explicitly represent the relations between any two nodes via edges (*e.g.*, foreign key relations). However, compared to

Approach	Spider-DK		Spider-Syn		Spider-Realistic	
	EM	EX	EM	EX	EM	EX
RAT-SQL + BERT (Wang et al. 2020a)	40.9	-	48.2	-	58.1	62.1
RAT-SQL + GRAPPA (Yu et al. 2021)	38.5	-	49.1	-	59.3	-
T5-3B (Gao et al. 2022)	-	-	59.4	65.3	63.2	65.0
LGESQL + ELECTRA (Cao et al. 2021)	48.4	-	64.6	-	69.2	-
TKK-3B (Gao et al. 2022)	-	-	63.0	68.2	68.5	71.1
T5-3B + PICARD (Qi et al. 2022)	-	-	-	-	68.7	71.4
RASAT + PICARD (Qi et al. 2022)	-	-	-	-	69.7	71.9
LGESQL + ELECTRA + SUN (Qin et al. 2022)	52.7	-	66.9	-	70.9	-
RESDSQL-3B + NatSQL	53.3	66.0	69.1	76.9	77.4	81.9

Table 3: Evaluation results on Spider-DK, Spider-Syn, and Spider-Realistic (%).

Model variant	Table AUC	Column AUC	Total
Cross-encoder	0.9973	0.9957	1.9930
- w/o enh. layer	0.9965	0.9939	1.9904
- w/o focal loss	0.9958	0.9943	1.9901

Table 4: Ablation studies of the cross-encoder.

Model variant	EM (%)	EX (%)
RESDSQL-Base	71.7	77.9
- w/o ranking schema items	67.2	70.1
- w/o skeleton parsing	71.0	77.1

Table 5: The effect of key designs.

PLMs, graph neural networks (GNNs) usually cannot be designed too deep due to the limitation of the over-smoothing issue (Chen et al. 2020), which restricts the representation ability of GNNs. Then, PLMs have already encoded language patterns in their parameters after pre-training (Zhang et al. 2021), however, the parameters of GNNs are usually randomized. Moreover, the graph encoder relies heavily on the design of relations, which may limit its robustness and generality on other datasets (Gao et al. 2022).

Grammar-Based Decoder To inject the SQL grammar into the decoder, Yin and Neubig (2017); Krishnamurthy, Dasigi, and Gardner (2017) propose a top-down decoder to generate a sequence of pre-defined actions that can describe the grammar tree of the SQL query. Rubin and Berant (2021) devise a bottom-up decoder instead of the top-down paradigm. PICARD (Scholak, Schucher, and Bahdanau 2021) incorporates an incremental parser into the auto-regressive decoder of PLMs to prune the invalid partially generated SQL queries during beam search.

Execution-Guided Decoder Some works use an off-the-shelf SQL executor such as SQLite to ensure grammatical correctness. Wang et al. (2018) leverage a SQL executor to check and discard the partially generated SQL queries which raise errors during decoding. To avoid modifying the decoder, Suhr et al. (2020) check the executability of each candidate SQL query, which is also adopted by our method.

Schema Item Classification

Schema item classification is often introduced as an auxiliary task to improve the schema linking performance for Text-to-SQL. For example, GRAPPA (Yu et al. 2021) and GAP (Shi et al. 2021) further pre-train the PLMs by using the schema item classification task as one of the pre-training objectives. Then, Text-to-SQL can be viewed as a downstream task to be fine-tuned. Cao et al. (2021) combine the schema item classification task with the Text-to-SQL task in a multi-task learning way. The above-mentioned methods enhance the encoder by pre-training or the multi-task learning paradigm. Instead, we propose an independent cross-encoder as the schema item classifier which is easier to be trained. We use the classifier to re-organize the input of the seq2seq model, which can produce a more direct impact on schema linking. Bogin, Gardner, and Berant (2019) calculate a relevance score for each schema item, which is then used as the soft coefficient of the schema items in the subsequent graph encoder. Compared with them, our method can be viewed as a hard filtering of schema items which can reduce noise more effectively.

Intermediate Representation

Because there is a huge gap between natural language questions and their corresponding SQL queries, some works have focused on how to design an efficient intermediate representation (IR) to bridge the aforementioned gap (Yu et al. 2018b; Guo et al. 2019; Gan et al. 2021b). Instead of directly generating full-fledged SQL queries, these IR-based methods encourage models to generate IRs, which can be translated to SQL queries via a non-trainable transpiler.

Conclusion

In this paper, we propose RESDSQL, a simple yet powerful Text-to-SQL parser. We first train a cross-encoder to rank and filter schema items which are then injected into the encoder of the seq2seq model. We also let the decoder generate the SQL skeleton first, which can implicitly guide the subsequent SQL generation. To a certain extent, such a framework decouples schema linking and skeleton parsing, which can alleviate the difficulty of Text-to-SQL. Extensive experiments on Spider and its three variants demonstrate the performance and robustness of RESDSQL.

Acknowledgments

We thank Hongjin Su and Tao Yu for their efforts in evaluating our model on Spider’s test set. We also thank the anonymous reviewers for their helpful suggestions. This work is supported by National Natural Science Foundation of China (62076245, 62072460, 62172424, 62276270) and Beijing Natural Science Foundation (4212022).

References

- Bogin, B.; Gardner, M.; and Berant, J. 2019. Global Reasoning over Database Structures for Text-to-SQL Parsing. In *EMNLP-IJCNLP 2019*, 3657–3662.
- Cai, R.; Xu, B.; Zhang, Z.; Yang, X.; Li, Z.; and Liang, Z. 2018. An Encoder-Decoder Framework Translating Natural Language to Database Queries. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, 3977–3983.
- Cai, R.; Yuan, J.; Xu, B.; and Hao, Z. 2021. SADGA: Structure-Aware Dual Graph Aggregation Network for Text-to-SQL. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021*, 7664–7676.
- Cao, R.; Chen, L.; Chen, Z.; Zhao, Y.; Zhu, S.; and Yu, K. 2021. LGESQL: Line Graph Enhanced Text-to-SQL Model with Mixed Local and Non-Local Relations. In *ACL/IJCNLP 2021*, 2541–2555.
- Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020. Measuring and Relieving the Over-Smoothing Problem for Graph Neural Networks from the Topological View. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 3438–3445.
- Deng, X.; Awadallah, A. H.; Meek, C.; Polozov, O.; Sun, H.; and Richardson, M. 2021. Structure-Grounded Pretraining for Text-to-SQL. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, 1337–1350.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019*, 4171–4186.
- Gan, Y.; Chen, X.; Huang, Q.; Purver, M.; Woodward, J. R.; Xie, J.; and Huang, P. 2021a. Towards Robustness of Text-to-SQL Models against Synonym Substitution. In *ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 2505–2515.
- Gan, Y.; Chen, X.; and Purver, M. 2021. Exploring Underexplored Limitations of Cross-Domain Text-to-SQL Generalization. In *EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, 8926–8931.
- Gan, Y.; Chen, X.; Xie, J.; Purver, M.; Woodward, J. R.; Drake, J. H.; and Zhang, Q. 2021b. Natural SQL: Making SQL Easier to Infer from Natural Language Specifications. In *Findings of EMNLP 2021*, 2030–2042.
- Gao, C.; Li, B.; Zhang, W.; Lam, W.; Li, B.; Huang, F.; Si, L.; and Li, Y. 2022. Towards Generalizable and Robust Text-to-SQL Parsing. In *Findings of EMNLP 2022*.
- Giordani, A.; and Moschitti, A. 2012. Automatic Generation and Reranking of SQL-derived Answers to NL Questions. In *Proceedings of the Second International Conference on Trustworthy External Systems via Evolving Software, Data and Knowledge*, 59–76.
- Guo, J.; Zhan, Z.; Gao, Y.; Xiao, Y.; Lou, J.; Liu, T.; and Zhang, D. 2019. Towards Complex Text-to-SQL in Cross-Domain Database with Intermediate Representation. In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28-August 2, 2019, Volume 1: Long Papers*, 4524–4535.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.*, 1735–1780.
- Hui, B.; Geng, R.; Wang, L.; Qin, B.; Li, Y.; Li, B.; Sun, J.; and Li, Y. 2022. S²SQL: Injecting Syntax to Question-Schema Interaction Graph Encoder for Text-to-SQL Parsers. In *Findings of ACL 2022*, 1254–1262.
- Iyer, S.; Konstas, I.; Cheung, A.; Krishnamurthy, J.; and Zettlemoyer, L. 2017. Learning a Neural Semantic Parser from User Feedback. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, 963–973.
- Krishnamurthy, J.; Dasigi, P.; and Gardner, M. 2017. Neural Semantic Parsing with Type Constraints for Semi-Structured Tables. In *EMNLP 2017*, 1516–1526.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *ACL 2020*, 7871–7880.
- Lin, T.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. In *IEEE International Conference on Computer Vision, ICCV 2017*, 2999–3007.
- Lin, X. V.; Socher, R.; and Xiong, C. 2020. Bridging Textual and Tabular Data for Cross-Domain Text-to-SQL Semantic Parsing. In *Findings of EMNLP 2020*, 4870–4888.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *7th International Conference on Learning Representations, ICLR 2019*.
- Qi, J.; Tang, J.; He, Z.; Wan, X.; Cheng, Y.; Zhou, C.; Wang, X.; Zhang, Q.; and Lin, Z. 2022. RASAT: Integrating Relational Structures into Pretrained Seq2Seq Model for Text-to-SQL. In *EMNLP 2022*.

- Qin, B.; Wang, L.; Hui, B.; Li, B.; Wei, X.; Li, B.; Huang, F.; Si, L.; Yang, M.; and Li, Y. 2022. SUN: Exploring Intrinsic Uncertainties in Text-to-SQL Parsers. In *Proceedings of the 29th International Conference on Computational Linguistics, COLING 2022, Gyeongju, Republic of Korea, October 12-17, 2022*, 5298–5308.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *J. Mach. Learn. Res.*, 140:1–140:67.
- Rubin, O.; and Berant, J. 2021. SmBoP: Semi-autoregressive Bottom-up Semantic Parsing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021*, 311–324.
- Schlichtkrull, M. S.; Kipf, T. N.; Bloem, P.; van den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling Relational Data with Graph Convolutional Networks. In *The Semantic Web - 15th International Conference, ESWC 2018*, 593–607.
- Scholak, T.; Schucher, N.; and Bahdanau, D. 2021. PICARD: Parsing Incrementally for Constrained Autoregressive Decoding from Language Models. In *EMNLP 2021*, 9895–9901.
- Shaw, P.; Chang, M.; Pasupat, P.; and Toutanova, K. 2021. Compositional Generalization and Natural Language Variation: Can a Semantic Parsing Approach Handle Both? In *ACL/IJCNLP 2021*, 922–938.
- Shaw, P.; Uszkoreit, J.; and Vaswani, A. 2018. Self-Attention with Relative Position Representations. In *NAACL-HLT*, 464–468.
- Shazeer, N.; and Stern, M. 2018. Adafactor: Adaptive Learning Rates with Sublinear Memory Cost. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, 4603–4611.
- Shi, P.; Ng, P.; Wang, Z.; Zhu, H.; Li, A. H.; Wang, J.; dos Santos, C. N.; and Xiang, B. 2021. Learning Contextual Representations for Semantic Parsing with Generation-Augmented Pre-Training. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021*, 13806–13814.
- Suhr, A.; Chang, M.; Shaw, P.; and Lee, K. 2020. Exploring Unexplored Generalization Challenges for Cross-Database Semantic Parsing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, 8372–8388.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, 5998–6008.
- Wang, B.; Shin, R.; Liu, X.; Polozov, O.; and Richardson, M. 2020a. RAT-SQL: Relation-Aware Schema Encoding and Linking for Text-to-SQL Parsers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, 7567–7578.
- Wang, C.; Tatwawadi, K.; Brockschmidt, M.; Huang, P.-S.; Mao, Y.; Polozov, O.; and Singh, R. 2018. Robust Text-to-SQL Generation with Execution-Guided Decoding. *arXiv preprint arXiv:1807.03100*.
- Wang, K.; Shen, W.; Yang, Y.; Quan, X.; and Wang, R. 2020b. Relational Graph Attention Network for Aspect-based Sentiment Analysis. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020*, 3229–3238.
- Xu, P.; Kumar, D.; Yang, W.; Zi, W.; Tang, K.; Huang, C.; Cheung, J. C. K.; Prince, S. J. D.; and Cao, Y. 2021. Optimizing Deeper Transformers on Small Datasets. In *ACL/IJCNLP 2021*, 2089–2102.
- Yin, P.; and Neubig, G. 2017. A Syntactic Neural Model for General-Purpose Code Generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017*, 440–450.
- Yu, T.; Li, Z.; Zhang, Z.; Zhang, R.; and Radev, D. R. 2018a. TypeSQL: Knowledge-Based Type-Aware Neural Text-to-SQL Generation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT*, 588–594.
- Yu, T.; Wu, C.; Lin, X. V.; Wang, B.; Tan, Y. C.; Yang, X.; Radev, D. R.; Socher, R.; and Xiong, C. 2021. GraPPa: Grammar-Augmented Pre-Training for Table Semantic Parsing. In *9th International Conference on Learning Representations, ICLR 2021*.
- Yu, T.; Yasunaga, M.; Yang, K.; Zhang, R.; Wang, D.; Li, Z.; and Radev, D. R. 2018b. SyntaxSQLNet: Syntax Tree Networks for Complex and Cross-Domain Text-to-SQL Task. *arXiv preprint arXiv:1810.05237*.
- Yu, T.; Zhang, R.; Yang, K.; Yasunaga, M.; Wang, D.; Li, Z.; Ma, J.; Li, I.; Yao, Q.; Roman, S.; Zhang, Z.; and Radev, D. R. 2018c. Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3911–3921.
- Zelle, J. M.; and Mooney, R. J. 1996. Learning to Parse Database Queries Using Inductive Logic Programming. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference, AAAI 96, IAAI 96*, 1050–1055.
- Zhang, Y.; Warstadt, A.; Li, X.; and Bowman, S. R. 2021. When Do You Need Billions of Words of Pretraining Data? In *ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, 1112–1125.
- Zhong, R.; Yu, T.; and Klein, D. 2020. Semantic Evaluation for Text-to-SQL with Distilled Test Suites. In *EMNLP 2020*, 396–411.
- Zhong, V.; Xiong, C.; and Socher, R. 2017. Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning. *arXiv preprint arXiv:1709.00103*.