Structure-Grounded Pretraining for Text-to-SQL

Xiang Deng [*1], Ahmed Hassan Awadallah[2], Christopher Meek[2],
Oleksandr Polozov[2], Huan Sun[1], and Matthew Richardson[2]

[1]The Ohio State University
{deng.595,sun.397}@osu.edu
[2]Microsoft Research, Redmond
{hassanam,meek,polozov,mattri}@microsoft.com

Abstract

Learning to capture text-table alignment is essential for tasks like text-to-SQL. A model needs to correctly recognize natural language references to columns and values and to ground them in the given database schema. In this paper, we present a novel weakly supervised **Stru**cture-**G**rounded pretraining framework (STRUG) for text-to-SQL that can effectively learn to capture text-table alignment based on a parallel text-table corpus. We identify a set of novel pretraining tasks: column grounding, value grounding and column-value mapping, and leverage them to pretrain a text-table encoder. Additionally, to evaluate different methods under more realistic text-table alignment settings, we create a new evaluation set Spider-Realistic based on Spider dev set with explicit mentions of column names removed, and adopt eight existing text-to-SQL datasets for cross-database evaluation. STRUG brings significant improvement over BERT$_{\text{LARGE}}$ in all settings. Compared with existing pretraining methods such as GRAPPA, STRUG achieves similar performance on Spider, and outperforms all baselines on more realistic sets. All the code and data used in this work is public available at https://aka.ms/strug.

1 Introduction

Semantic parsing is the task of mapping a natural language (NL) utterance to a machine-understandable representation such as lambda calculus, abstract meaning representation, or a structured query language (e.g., SQL). In this paper, we focus on the task of translating NL questions to executable SQL queries (text-to-SQL). This is a fundamental task for building natural language interfaces for databases, which can enable non-expert users to effortlessly query databases (Androutsopoulos et al., 1995; Li and Jagadish, 2014a).
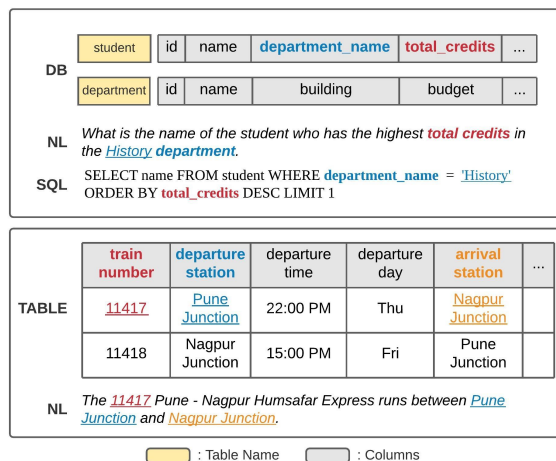

Figure 1: Illustration of text-to-SQL text-table alignment (top half) and parallel text-table corpus (bottom half). In both examples, the associations between tokens in the NL utterance and columns in the table are indicated. In this paper, we aim to leverage the text-table alignment knowledge in the parallel text-table corpus to help text-to-SQL.

One of the key challenges in text-to-SQL is text-table alignment, that is, *to correctly recognize natural language references to columns and values and to ground them in the given database schema*. Consider the example in the top half of Fig. 1. A model needs to first identify the column mentions `total credits`, `department`, and value mention `History`, and then ground them to the given schema. This is challenging for three reasons. First, the model needs to jointly understand the NL utterance and the database schema, as the user may refer to a column using various expressions which usually differ from the original column name. Second, the model needs to be able to generalize to new database schemas and referential language that is not seen in training. Finally, in the case that accessing cell values is not possible, the model still needs to identify potential value mentions and link them to the correct columns without exhaustively searching and matching over the database.

1337

On the other hand, text-table alignment naturally exists in parallel text-table corpora, e.g., web tables with context (Lehmberg et al., 2016), table-to-text generation datasets (Parikh et al., 2020; Chen et al., 2020a), table-based question answering datasets (Pasupat and Liang, 2015; Chen et al., 2020b). Such datasets can be collected from web pages, documents, etc., and requires much less human effort to create compared with text-to-SQL datasets. The bottom half of Fig. 1 gives an example of such an alignment dataset. There are three value mentions `11417`, `Pune Junction` and `Nagpur Jnction`, which can be grounded to the `train number`, `departure station` and `arrival station` columns respectively. Such alignment information can be easily obtained by leveraging the table contents or using some human annotation. In this work, we aim to incorporate the text-table alignment knowledge contained in a parallel corpus via pretraining and use it to help the downstream text-to-SQL task.

We present a novel weakly supervised structure-grounded pretraining framework (STRUG) for text-to-SQL. We design a set of prediction tasks and optimize them leveraging a parallel corpus containing both NL sentences and tabular data to encourage the encoded representation to capture information required to support tasks that require table grounding. More specifically, we identify three critical tasks for aligning text with table: column grounding, value grounding and column-value mapping (examples shown in Fig. 2). We re-purpose an existing large-scale table-to-text generation dataset ToTTo (Parikh et al., 2020) for pretraining and gain labels for the three tasks via weak supervision. We experiment under two settings, with or without human assistance: (1) *human assisted setting*, using ToTTo's revised descriptions and cell annotations; (2) *automatic setting*, using the raw sentences and inferring the cell correspondences via string matching with the table contents.

As pointed out by Suhr et al. (2020), existing text-to-SQL benchmarks like Spider (Yu et al., 2018b) render the text-table alignment challenge easier than expected by explicitly mentioning exact column names in the NL utterances. Contrast this to more realistic settings where users may refer to the columns using a variety of expressions. Suhr et al. (2020) propose a new cross-database setting that uses Spider for training and includes eight other single-domain text-to-SQL datasets for
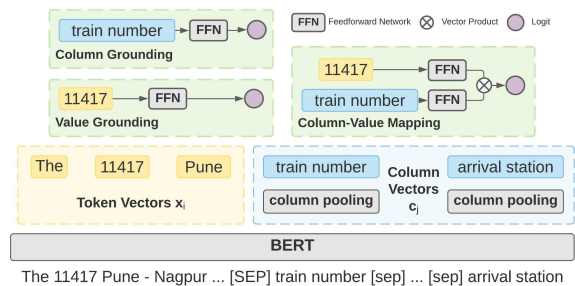

Figure 2: Overview of our model architecture and three pretraining objectives.

evaluation. In addition to adopting their setting, we create a new evaluation set called Spider-Realistic from the original Spider dev set, by removing explicit mentions of column names from an utterance.

We pretrain STRUG using 120k text-table pairs from ToTTo. Experiments show that our structure-grounded pretraining objectives are very efficient and usually converge with around 5 epochs in less than 4 hours. This dramatically reduces the pretraining cost compared to previous pretraining methods (Herzig et al., 2020; Yin et al., 2020). We adopt the same model architecture as BERT (Devlin et al., 2019), with simple classification layers on top for pretraining. For downstream tasks, STRUG can be used as a text-table encoder and easily integrated with any existing state-of-the-art model. We conduct extensive experiments and show that:

(1) Combined with state-of-the-art text-to-SQL model RAT-SQL (Wang et al., 2020), using STRUG as encoder significantly outperforms directly adopting pretrained BERT$_{LARGE}$ (RAT-SQL's default encoder) and performs on par with other text-table pretraining models like GRAPPA (Yu et al., 2020) on the widely used Spider benchmark.

(2) On more realistic evaluation settings, including Spider-Realistic and the Suhr et al. (2020) datasets, our method outperforms all baselines. This demonstrates the superiority of our pretraining framework in solving the text-table alignment challenge, and its usefulness in practice.

(3) STRUG also helps reduce the need for large amount of costly supervised training data. We experiment with the WikiSQL benchmark (Zhong et al., 2017) by limiting training data size, and show that our pretraining method can boost the model performance by a large margin and consistently outperforms existing pretraining methods.

2 Related Work

Cross-Database Text-to-SQL. Remarkable progress has been made in text-to-SQL over the past few years. With sufficient in-domain training data, existing models already achieve over 80% exact matching accuracy (Finegan-Dollak et al., 2018; Wang et al., 2018) on single-domain benchmarks like ATIS (Hemphill et al., 1990; Dahl et al., 1994) and GeoQuery (Zelle and Mooney, 1996). However, annotating NL questions with SQL queries is expensive making it cost-prohibitive to collect training examples for all possible databases. A model that can generalize across domains and databases is desired. In light of this, Yu et al. (2018b) present Spider, a cross-database text-to-SQL benchmark that trains and evaluates a system using different databases. More recently, Suhr et al. (2020) provide a holistic analysis of the challenges introduced in cross-database text-to-SQL and propose to include single-domain datasets in evaluation. Their study uncovers the limitations of current text-to-SQL models, and demonstrates the need for models that can better handle the generalization challenges.

Pretraining for Text-Table Data. Inspired by the success of pretrained language models, some recent work has tried to apply similar pretraining objectives to text-table data. TaBERT (Yin et al., 2020) and TAPAS (Herzig et al., 2020) jointly learn text-table representations by leveraging a large amount of web tables and their textual context. They flatten the tables and use special embeddings to model the structure information. A masked language model (MLM) objective is then used to predict the masked tokens in the text-table data. MLM is good at modeling the contextualized semantic representations of a token, but is weak at capturing the alignment between a pair of sequences (e.g., text-table). More recently, GRAPPA (Yu et al., 2020) explores a different direction for pretraining which shares some similarity with existing work on data augmentation for semantic parsing. GRAPPA first constructs synthetic question-SQL pairs using templates (a synchronous context free grammar) induced from existing text-to-SQL datasets, a SQL semantic prediction objective is then used to learn compositional inductive bias from the synthetic data. However, as the synthetic data is generated using templates, and the column names and values are directly filled in the questions, it has the same problem as existing text-to-SQL datasets that

eases the text-table alignment challenge. In contrast, STRUG aims to directly learn the text-table alignment knowledge from parallel text-table corpora via structure-grounded pretraining objectives. We also note that existing pretraining methods and STRUG can be complementary and combined together in the future.

Structure Grounding in Text-to-SQL. Structure grounding has been proven to be crucial for text-to-SQL, where a model needs to correctly identify column and value mentions in an NL utterance and link them to the given database schema (Guo et al., 2019; Bogin et al., 2019; Wang et al., 2020; Lei et al., 2020). Most existing text-to-SQL systems have specially designed components for structure grounding, which is also referred to as schema linking. For example, Guo et al. (2019); Yu et al. (2018a) explore using simple heuristics like string matching for schema linking, and use the linking results as direct hints to their systems. However, such heuristics may not generalize well in real world scenarios where there are varied ways to refer to a column, which usually differ from the original column name. More recently, Shi et al. (2020) and Lei et al. (2020) take a step forward and manually annotate WikiTableQuestions (Pasupat and Liang, 2015) and Spider with fine-grained alignment labels for supervised training (together with the text-to-SQL objective), which brings significant improvements. The main drawback of these models is that they are limited to learn the alignment knowledge from a relatively small training corpus, and cannot generalize well in a cross-domain setting. Moreover, SQL annotations and fine-grained alignment labels are both expensive to get manually. In contrast, this paper aims to re-purpose an existing parallel text-table corpus for pretraining models to learn structure grounding, where we generate alignment labels at large scale with low or no cost.

3 Structure-Grounded Pretraining

3.1 Motivation

One of the critical generalization challenges in cross-database text-to-SQL is text-table alignment, i.e., a model needs to understand NL utterances and database schemas unseen in training, including value mentions and novel columns, and to correctly map between them. Similar generalization challenges have been studied for a long time in the NLP field. Recently, pretrained language models (Devlin et al., 2019; Liu et al., 2019; Lewis et al.,

Year	Competition	Venue	Position	Event	Notes
1993	European Junior Championships	San Sebastian, Spain	7th	100m	11.74
			3rd	4x100 m relay	44.60
1994	World Junior Championships	Lisbon, Portugal	12th (semis)	100m	11.66 (wind: +1.3 m/s)
			2nd	4x100 m relay	44.78
1995	World Championships	Gothenburg, Sweden	7th (q-finals)	100m	11.54
			3rd	4x100 m relay	43.01

	Human Assisted Setting	**Automatic Setting**
NL Utterance	Gabriele Becker competed at the 1995 World Championships both individually and in the relay.	After winning the German under-23 100 m title, she was selected to run at the 1995 World Championships in Athletics both individually and in the relay.
Column Grounding	Year, Competition, Event	Year, Competition
Value Grounding	1995, World Championships	1995, World Championships
Column-Value Maping	1995-Year, World Championships-Competition	1995-Year, World Championships-Competition

Figure 3: Illustration of the parallel corpus ToTTo (Parikh et al., 2020) and our two weakly supervised pretraining settings. Cell highlighted with yellow are the cell annotations provided by ToTTo, and cell highlighted with dashed lines are cell annotations obtained via string matching in automatic setting.

2020) have achieved great success in tackling the challenges by learning contextualized representations of words from a large text corpus. Inspired by this, in this work we aim to develop a pretraining method that can directly learn the text-table alignment knowledge from a large parallel text-table corpus.

Unlike previous text-table pretraining works (Herzig et al., 2020; Yin et al., 2020) that optimize unsupervised objectives like MLM during pretraining, we carefully design three structure-grounded tasks: column grounding, value grounding and column-value mapping. These tasks are related to text-to-SQL and can directly capture the text-table alignment during pretraining. As a result, the learned alignment knowledge can be effectively transferred to the downstream task and improve the final performance.

3.2 Pretraining Objectives

We use the same model architecture as BERT, and add simple classification layers on top for the three structure-grounded tasks. For downstream tasks, our model can be easily integrated into existing models as text-table encoder. Following previous work (Hwang et al., 2019; Wang et al., 2020; Guo et al., 2019), we linearize the input by concatenating the NL utterance and column headers, using `<sep>` token as a separator.

Formally, given a pair of NL utterance $\{x_i\}$ and table with a list of column headers (in case there are multiple tables like in databases, we concatenate all the column names together) $\{c_j\}$, we first obtain the contextualized representation $\mathbf{x}_i$ of each token in the utterance and $\mathbf{c}_j$ for each column using the last layer output of the BERT encoder. Here each column header c_j may contain multiple tokens $c_{j,0}, \ldots, c_{j,|c_j|}$. We obtain a single vector representation for each column using column pooling. More specifically, we take the output of the first and last token of the header, and calculate the column representation as $\mathbf{c}_j = (\mathbf{c}_{j,0} + \mathbf{c}_{j,|c_j|})/2$. $\{\mathbf{x}_i\}$ and $\{\mathbf{c}_j\}$ are then used to compute losses for the three tasks. An overview of our model architecture and pretraining objectives are shown in Fig. 2.

Column grounding. An important task in text-to-SQL is to identify grounded columns from the schema and use them for the generated SQL query. With a parallel text-table corpus, this is similar to selecting the columns that are mentioned in the associated NL sentence. This task requires a model to understand the semantic meaning of a column based on its header alone, and to infer its relation with the NL sentence based on the contextualized representations. We formulate it as a binary classification task. For each column c_j, we use a one-layer feed forward network $f(\cdot)$ to get prediction $p_j^c = f(\mathbf{c}_j)$ of whether c_j is mentioned in the sentence or not. The column grounding loss $\mathcal{L}_c$ is then calculated using the binary cross entropy loss w.r.t. ground truth labels $y_j^c \in \{0, 1\}$. Note this task requires the model to identify the meaning of a column without access to any of its values. Hence, it is suitable for the typical text-to-SQL setting where the model only has access to the database schema.

Value grounding. For clauses like `WHERE` and `HAVING`, to generate an executable SQL query, a model also needs to extract the value to be compared with the grounded column from the NL utterance. This can be transformed to the task of finding cell mentions in the NL sentence with a parallel text-table corpus. Since the contents of the table is

Dataset	# Examples	Exec Acc (Suhr et al., 2020)	% Col Mentioned
ATIS (Hemphill et al., 1990; Dahl et al., 1994)	289 (486)	0.8	0.0
Restaurants (Tang and Mooney, 2000)	27 (378)	3.7	0.0
Academic(Li and Jagadish, 2014b)	180 (196)	8.2	11.4
Yelp(Yaghmazadeh et al., 2017)	54 (128)	19.8	8.0
Scholar(Iyer et al., 2017)	394 (599)	0.5	0.0
Advising(Finegan-Dollak et al., 2018)	309 (2858)	2.3	0.3
IMDB(Yaghmazadeh et al., 2017)	107 (131)	24.6	1.0
GeoQuery(Zelle and Mooney, 1996)	532 (598)	41.6	3.9
Spider (Yu et al., 2018b)	1034	69.0	39.2
Spider-Realistic	508	-	1.8

Table 1: Statistic of the datasets used in this work. Here we show the number of examples for evaluation after filtering (sizes of the original datasets before any filtering are shown in parentheses), and the execution accuracy reported in Suhr et al. (2020). For the detailed filtering process of Suhr et al. (2020), please check the original paper or Appendix A.1. % Col Mentioned[1] measures the proportion of examples in the evaluation set where all columns compared against entities in the gold query are explicitly mentioned in the NL utterance.

not available, it is necessary for the model to infer the possible value mentions based on NL utterance and the table schema only. Similarly to column grounding, we also view this as a classification task. For each token x_i, we get prediction of x_i being part of a grounded value as $p_i^v = f(\mathbf{x}_i)$. The value grounding loss $\mathcal{L}_v$ is then calculated using the binary cross entropy loss w.r.t. ground truth labels $y_i^v \in \{0, 1\}$.

Column-Value mapping. As there may be multiple columns and values used in the SQL query, a text-to-SQL model also needs to correctly map the grounded columns and values. This is used to further strengthen the model's ability to capture the correlation between the two input sequences by learning to align the columns and values. We formulate this as a matching task between the tokens in the NL sentence and the columns. For every grounded token x_i (i.e., $y_i^v = 1$), we pair it with each column c_j and calculate the probability of x_i matching c_j as $p_{i,j}^{cv} = f([\mathbf{x}_i, \mathbf{c}_j])$. Here $[\cdot, \cdot]$ is the vector concatenation operation. We then apply a softmax layer over the predictions for each token $p_i^{cv} = \{p_{i,j}^{cv}\}_{j=1}^{|c|}$, and the final column-value mapping loss $\mathcal{L}_{cv}$ is then calculated as $\mathcal{L}_{cv} = \text{CrossEntropy}(\text{softmax}(p_i^{cv}), y_i^{cv})$, where $y_i^{cv} \in \{0, 1\}^{|c|}$ is the ground truth label.

The final loss $\mathcal{L}$ for pretraining is the sum of all three losses. We experimented with different weights for each term, but did not observe significant improvement on the results. Hence we only report results with equally weighted losses.

$$\mathcal{L} = \mathcal{L}_c + \mathcal{L}_v + \mathcal{L}_{cv} \qquad (1)$$

3.3 Obtaining Pretraining Data via Weak Supervision

We obtain ground truth labels y_j^c, y_i^v and y_i^{cv} from a parallel text-table corpus based on a simple intuition: given a column in the table, if any of its cell values can be matched to a phrase in the sentence, this column is likely mentioned in the sentence, and the matched phrase is the value aligned with the column. To ensure high quality text-table alignment information in the pretraining corpus, unlike previous work (Herzig et al., 2020; Yin et al., 2020) that use loosely connected web tables and their surrounding text, here we leverage an existing large-scale table-to-text generation dataset ToTTo (Parikh et al., 2020). ToTTo contains 120,761 NL descriptions and corresponding web tables automatically collected from Wikipedia using heuristics. Additionally, it provides cell level annotation that highlights cells mentioned in the description and revised version of the NL descriptions with irrelevant or ambiguous phrases removed.

We experiment with two pretraining settings, with or without human assistance. In the *human assisted setting*, we use the cell annotations along with the revised description to infer the ground truth labels. More specifically, we first label all the columns c_j that contain at least one highlighted cell as positive ($y_j^c = 1$). We then iterate through all the values of the highlighted cells and match them with the NL description via exact string matching to extract value mentions. If a phrase is matched to a highlighted cell, we select all the tokens x_i in that phrase and align them with the corresponding

---

[1]Unlike Suhr et al. (2020), here we do not consider examples where there is no column compared against entity.

columns c_j ($y_i^v = 1$, $y_{i,j}^{cv} = 1$). In the *automatic setting*, we use only the tables and the raw sentences, and obtain cell annotations by comparing each cell with the NL sentence using exact string matching. Note that in both settings, the cell values are used only for preparing supervision for the pretraining objectives, not as inputs to the pretraining model.

To make the pretraining more effective and to achieve a better generalization performance, we also incorporate two data augmentation techniques. First, since the original parallel corpus only contains one table for each training example, we randomly sample K_{neg} tables as negative samples and append their column names to the input sequence. This simulates a database with multiple tables and potentially hundreds of columns, which is common in text-to-SQL. Second, we randomly replace the matched phrases in the NL sentences with values of cells from the same column (the labels are kept the same). This way we can better leverage the contents of the table during pretraining and improve the model's generalization ability by exposing it to more cell values.

4 Creating a More Realistic Evaluation Set

As one of the first datasets to study cross-database text-to-SQL, Spider has been a widely used benchmark in assessing a model's ability to generalize to unseen programs and databases. However, as pointed out by Suhr et al. (2020), Spider eases the task by using utterances that closely match their paired SQL queries, for example by explicitly mentioning the column names in the question, while in practice NL references to columns usually differ from the original column name. To alleviate this problem, Suhr et al. (2020) propose to train the model with cross-domain dataset like Spider, and add another eight single-domain datasets like ATIS (Hemphill et al., 1990; Dahl et al., 1994) and Geo-Query (Zelle and Mooney, 1996) for evaluation. However, some of the datasets differ a lot from Spider, introducing many novel query structures and dataset conventions.[2] As we can see from Table 1, their model (Suhr et al., 2020) has very poor performance in some datasets. In light of this, we present a new realistic and challenging evaluation set based on Spider. We first select a complex subset from

Example	Type
Show name, country, age for all singers ~~ordered by *age*~~ from the oldest to the youngest.	Remove
Find the number of concerts happened in the stadium ~~with the highest *capacity*~~ that can accommodate the most people. How many pets ~~have a greater *weight* than 10~~ are over 10 lbs?	paraphrase

Table 2: Examples of how we create Spider-Realistic from Spider. Phrases shown in italic exactly match with column names.

the Spider dev set where there are columns compared against values or used in clauses like `ORDER BY`. We then manually modify the NL questions in the subset ourselves to remove or paraphrase explicit mentions of columns names, except for the columns in `SELECT` clauses, while keeping the SQL queries unchanged. Some examples are shown in Table 2. This way we do not introduce extra challenges like adapting to new query structures but make it possible to fairly assess the model's capability in aligning text and tables. To make a more comprehensive comparison, we will also report results on the original Suhr et al. (2020) datasets.

5 Experiments

5.1 Benchmarks and Base Models

Spider and the realistic evaluation sets. Spider (Yu et al., 2018b) is a complex cross-database text-to-SQL dataset. It contains 10k complex question-query pairs grounded on 200 databases where multiple tables are joined via foreign keys. In addition, we create a new realistic evaluation set Spider-Realistic as described in Section 4. We also include the original Suhr et al. (2020) datasets, for a more comprehensive comparison. For the base model, we use RAT-SQL (Wang et al., 2020) which is the state-of-the-art model according to the official leaderboard as of the submission time. To generate executable SQL queries, we modify the pointer generator in RAT-SQL to enable it to copy values from the question. We use the same trained model for evaluation on the Spider dev set and the realistic evaluation sets. Yu et al. (2018b) includes some single-domain text-to-SQL datasets like GeoQuery as extra training data for Spider. Following Suhr et al. (2020), we train the model with only the original Spider data, and discard additional training data used by some previous works like Yu et al. (2018b). We use both the set match accuracy (exact match)

---

[2] Some of the datasets contain operators that are not covered by Spider grammar or novel query structure like self join that does not exist in the training corpus.

	Models	Spider-Realistic	ATIS	GeoQuery	Restaurants	Academic	IMDB	Yelp	Scholar	Advising
	# Examples	508	289	532	27	180	107	54	394	309
Schema Only	Suhr et al. (2020)	-	0.8 (0.5)	41.6 (35.6)	3.7 (3.7)	8.2 (6.1)	24.6 (24.3)	**19.8 (16.7)**	0.5 (0.4)	2.3 (1.2)
	RAT-SQL *w/o* value linking									
	w. BERT$_{LARGE}$	52.4 ± 0.7 (46.9)	2.1 ± 0.6	41.2 ± 11.6	0.0 ± 0.0	5.9 ± 2.1	26.5 ± 5.0	12.3 ± 1.7	0.8 ± 0.4	1.6 ± 0.4
	w. STRUG (Human Assisted)	57.8 ± 0.6 (53.3)	2.2 ± 0.2	45.5 ± 1.8	11.1 ± 9.1	**14.8 ± 5.0**	**37.1 ± 1.8**	15.4 ± 0.9	4.3 ± 1.7	**2.2 ± 0.4**
	w. STRUG (Automatic)	**60.3 ± 0.7 (54.9)**	**2.2 ± 0.2**	**50.9 ± 4.0**	**40.7 ± 5.2**	12.4 ± 1.9	35.5 ± 2.0	13.0 ± 2.6	**5.4 ± 0.7**	1.0 ± 0.3
Content Used	RAT-SQL									
	w. BERT$_{LARGE}$	62.1 ± 1.3 (58.1)	2.3 ± 0.2	47.3 ± 3.7	37.0 ± 18.9	15.6 ± 2.0	21.8 ± 1.6	16.0 ± 3.1	3.4 ± 1.4	6.4 ± 2.3
	w. GRAPPA	- (59.3)								
	w. STRUG (Human Assisted)	**65.7 ± 0.7 (62.2)**	**5.5 ± 1.1**	**59.5 ± 3.2**	40.7 ± 13.9	18.7 ± 2.1	26.8 ± 2.9	**21.6 ± 2.3**	**6.3 ± 1.8**	**6.9 ± 0.6**
	w. STRUG (Automatic)	65.3 ± 0.7 (62.2)	2.8 ± 0.7	57.5 ± 0.2	**44.4 ± 32.7**	**20.2 ± 1.6**	**30.2 ± 5.8**	18.5 ± 1.5	6.1 ± 0.5	5.2 ± 0.5

Table 3: Execution accuracy on the more realistic evaluation sets including Spider-Realistic and the Suhr et al. (2020) evaluation sets. For Spider-Realistic, we also show exact match accuracy in parentheses. For the Suhr et al. (2020) evaluation sets, we show results for the filtered set where examples with query returning empty set are excluded. Suhr et al. (2020) uses the WikiSQL dataset as additional training data, and we also show their results with only the Spider training data in parentheses.

	Models	Exact	Exec	Exact (Test)
Schema Only	EditSQL (Zhang et al., 2019) w. BERT	57.6	-	53.4
	IRNET (Guo et al., 2019) w. BERT	61.9	-	54.7
	RYANSQL (Choi et al., 2020) w. BERT	**70.6**	-	60.6
	Suhr et al. (2020) w. BERT$_{LARGE+}$	65.0	69.0	-
	RAT-SQL (Wang et al., 2020) *w/o* value linking			
	w. BERT$_{LARGE}$	67.0 ± 0.6	69.8 ± 0.3	-
	w. STRUG (Human Assisted)	70.5 ± 0.6	73.3 ± 0.4	**67.4**
	w. STRUG (Automatic)	69.8 ± 0.3	**74.2 ± 0.8**	-
Content Used	Global-GNN (Bogin et al., 2019)	52.7	-	47.4
	TranX w. TaBERT (Yin et al., 2020)	64.5	-	-
	RAT-SQL			
	w. BERT$_{LARGE}$	69.8 ± 0.8	72.3 ± 0.6	-
	w. GRAPPA (Yu et al., 2020)	**73.4**	-	**69.6**
	w. STRUG (Human Assisted)	72.7 ± 0.7	**75.5 ± 0.8**	68.4
	w. STRUG (Automatic)	72.6 ± 0.1	74.9 ± 0.1	-

Table 4: Results on Spider. The top half shows models using only database schema, the bottom half shows models using the database content. We train our model three times with different random seeds and report the mean and standard deviation here.

Models	ACC$_{lf}$	ACC$_{ex}$
HydraNet (Lyu et al., 2020)	**83.8**	**89.2**
X-SQL (He et al., 2019)	83.3	88.7
SQLova (Hwang et al., 2019)		
w. BERT$_{LARGE}$	82.1	87.3
w. TaBERT	**82.5**	**87.9**
w. STRUG (Human Assisted)	82.1	87.5
w. STRUG (Automatic)	82.4	87.8
SQLova (5%)		
w. BERT$_{LARGE}$	70.7	77.0
w. TaBERT	71.5	78.0
w. STRUG (Human Assisted)	**75.6**	**81.6**
w. STRUG (Automatic)	75.6	81.4

Table 5: Performance on WikiSQL. Here we show logical form accuracy and execution accuracy on the test set. (5%) means random sampling 5% of original training data for training.

from the official Spider evaluation script and execution accuracy[3] for evaluation on Spider and Spider-Realistic. On the Suhr et al. (2020) datasets, we use the official evaluation script[4] released by the authors and report execution accuracy.

WikiSQL. WikiSQL (Zhong et al., 2017) is a large-scale text-to-SQL dataset consists of over 80k question-query pairs grounded on over 30k Wikipedia tables. Although existing models are already reaching the upper-bound performance on this dataset (Hwang et al., 2019; Yavuz et al., 2018), mainly because of the simplicity of the SQL queries and large amount of data available for training, previous works have also used this dataset to demonstrate the model's generalization ability with limited training data (Yu et al., 2020; Yao et al., 2020). For the base model, we use SQLova (Hwang et al., 2019) without execution-guided decoding. Following the official leaderboard, we report both logical form accuracy and execution accuracy.

5.2 Training Details

For all experiments, we use the BERT implementation from Huggingface (Wolf et al., 2020) and the pretrained BERT$_{LARGE}$ model from Google [5]. For pretraining, we use Adam optimizer (Kingma and Ba, 2015) with a initial learning rate of 2e-5 and batch size of 48. In both settings, we use $K_{neg} = 1$ and pretrains our model for 5 epochs. We use 4 V100 GPUs for pretraining, which takes less than 4 hours.

For Spider and the realistic evaluation sets, we use the official implementation of RAT-SQL [6] and modify it to generate executable SQL queries. We follow the original settings and do hyperparameter search for learning rate (3e-4, 7.44e-4) and

---

[3]We execute case insensitive SQL queries, and compare the returned table.

[4]https://github.com/google-research/language/tree/master/language/xsp

---

[5]We use the BERT-Large, Uncased (Whole Word Masking) model from https://storage.googleapis.com/bert_models/2019_05_30/wwm_uncased_L-24_H-1024_A-16.zip

[6]https://github.com/microsoft/rat-sql

warmup step (5k, 10k). We use the same polynomial learning rate scheduler with warmup and train for 40,000 steps with batch size of 24. The learning rate for the pretrained encoder (e.g. BERT) is 3e-6 and is frozen during warmup.

For WikiSQL, we use the official SQLova implementation [7]. We use the default setting with learning rate of 1e-3 for the main model and learning rate of 1e-5 for the pretrained encoder. We train the model for up to 50 epochs and select the best model using the dev set.

5.3 Main Results

Spider. We first show results on Spider dev set in Table 4. The original Spider setting assumes only the schema information about the target database is known in both training and evaluation phase, as the content of the database may not be accessible to the system due to privacy concern. More recently, some works have tried to using the database content to help understand the columns and link with the NL utterance. Here we show results for both settings. In the first setting where only schema information is known, we disable the value-based linking module in RAT-SQL. As we can see from Table 4, replacing BERT$_{LARGE}$ with STRUG consistently improves the model performance in both settings. Under the setting where content is available, using STRUG achieves similar performance as GRAPPA and outperforms all other models. GRAPPA uses both synthetic data and larger text-table corpus for pretraining. However, it mainly learns inductive bias from the synthetic data while our model focuses on learning text-table association knowledge from the parallel text-table data. In error analysis on the Spider dev set, we notice that our best model[8] corrects 76 out of 270 wrong predictions made by GRAPPA while GRAPPA corrects 80 out of 274 wrong predictions made by our model. This demonstrates that the two pretraining techniques are complementary and we expect combining them can lead to further performance improvement. For results on different difficulty levels and components, please see Appendix B.1.

More realistic evaluation sets. Results on the realistic evaluation sets are summarized in Table 3. Firstly, we notice the performance of all models drops significantly on Spider-Realistic, demonstrating that inferring columns without explicit hint is
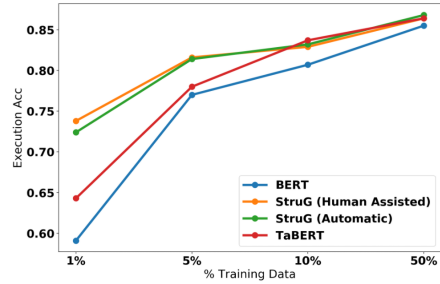

Figure 4: Execution Accuracy on the WikiSQL test set with different fractions of training data.
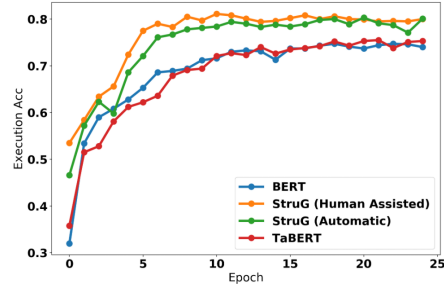

Figure 5: Execution Accuracy on the WikiSQL dev set during training with 5% of training data.

a challenging task and there is much room for improvement. Secondly, using STRUG brings consistent improvement over BERT$_{LARGE}$ in all realistic evaluation sets. In the Spider-Realistic set, using STRUG also outperforms GRAPPA[9] by 2.9%. Under the original Suhr et al. (2020) setting, combining RAT-SQL with STRUG significantly outperforms Suhr et al. (2020) in all datasets, despite that we do not include WikiSQL as additional training data as they did. Thirdly, comparing results in Table 4 with Table 3, using STRUG brings larger improvement over BERT$_{LARGE}$ in the more realistic evaluation sets. As shown in Table 1, the original Spider dataset has a high column mention ratio, so the models can use exact match for column grounding without really understanding the utterance and database schema. The more realistic evaluation sets better simulate the real world scenario and contain much less such explicit clues, making the text-table alignment knowledge learned by STRUG more valuable. For case studies on Spider-Realistic, please check Section 5.4.

WikiSQL. Results on WikiSQL are summarized in Table 5. When using the full training corpus, we notice that using STRUG achieves similar performance as BERT$_{LARGE}$. This is probably because of

---

[7] https://github.com/naver/sqlova
[8] RAT-SQL *w.* STRUG (Human Assisted)

[9] We use the checkpoint provided by the author, which achieves 73.8% exact match accuracy on the Spider dev set. Here we only evaluate on Spider-Realistic with exact match accuracy because their model does not generate values and includes IMDB and Geo as extra training data.

Spider-Realistic	What are the names of tournaments that have more than 10 matches?
	w. STRUG (Automatic) ✓
	`SELECT tourney_name FROM matches`
	`GROUP BY tourney_name`
	`HAVING Count(*) > 10`
	w. BERT_{LARGE} ✗
	`SELECT first_name FROM players JOIN`
	`matches GROUPBY first_name`
	`HAVING Count(*) > 10`

IMDB	List " James Bond " directors
	w. STRUG (Automatic) ✓
	`SELECT name FROM director`
	`JOIN directed_by JOIN MOVIE`
	`WHERE movie.title = "james bond"`
	w. BERT_{LARGE} ✗
	`SELECT gender FROM director`
	`WHERE director.name = "james bond"`

Table 6: Case study.

the large size of training data and the simple SQL structure of WikiSQL. To better demonstrate that the knowledge learned in pretraining can be effectively transferred to text-to-SQL task and reduce the need for supervised training data, we also conduct experiments with randomly sampled training examples. From Fig. 4 we can see that with only 1% of training data (around 500 examples), models using STRUG can achieve over 0.70 accuracy, outperforming both BERT_{LARGE} and TaBERT by a large margin. STRUG brings consist improvement over BERT_{LARGE} until we use half of the training data, where all models reach nearly the same performance as using the full training data. We also show the training progress using 5% of training data in Fig. 5. We can see that STRUG also helps speed up the training progress. For more break-down results on several subtasks, please see Appendix B.2.

Comparison of human assisted and automatic setting. In all benchmarks, we notice that STRUG pretrained using the automatic setting actually performs similarly as the setting where cell annotations are used. This indicates the effectiveness of our heuristic for cell annotation and the potential to pretrain STRUG with more unannotated parallel text-table data.

5.4 Case Study

We compare the predictions made by RAT-SQL *w.* BERT_{LARGE} and *w.* STRUG (Automatic). Some examples are shown in Table 6. In the first example from Spider-Realistic, we can see that the model *w.* BERT_{LARGE} fails to align *tournaments* with the *tourney_name* column, because of string mismatch.

In the second example from IMDB, although the model correctly recognizes *James Bond* as value reference, it fails to ground it to the correct column which is *movie_title*. This supports our hypothesis that using STRUG helps to improve the structure grounding ability of the model.

6 Conclusion

In this paper, we propose a novel and effective structure-grounded pretraining technique for text-to-SQL. Our approach to pretraining leverages a set of novel prediction tasks using a parallel text-table corpus to help solve the text-table alignment challenge in text-to-SQL. We design two settings to obtain pretraining labels without requiring complex SQL query annotation: using human labeled cell association, or leveraging the table contents. In both settings, STRUG significantly outperforms BERT_{LARGE} in all the evaluation sets. Meanwhile, although STRUG is surprisingly effective (using only 120k text-table pairs for pretraining) and performs on par with models like TaBERT (using 26m tables and their English contexts) and GRAPPA (using 475k synthetic examples and 391.5k examples from existing text-table datasets) on Spider, we believe it is complementary with these existing text-table pretraining methods. In the future, we plan to further increase the size of the pretraining corpus, and explore how to incorporate MLM and synthetic data.

Ethical Considerations

Dataset. In this work, we re-purpose an existing table-to-text generation dataset ToTTo (Parikh et al., 2020) for our pretraining. We obtain labels for our three pretraining tasks via weak supervision, which uses only the raw sentence-table pairs, or the cell annotations and revised descriptions that are already included in ToTTo dataset. As a result, no extra human effort is required for collecting our pretraining corpus. We also curate a more realistic evaluation dataset for text-to-SQL based on Spider dev set. In particular, we first select a complex subset from the Spider dev set and manually revise the NL questions to remove the explicit mention of column names. The detailed description of the process can be found in Section 4. The first author manually revised all the questions himself, which results in 508 examples in total.

Application. We focus on the task of text-to-SQL, which is a fundamental task for building natural language interfaces for databases. Such interface

can enable non-expert users to effortlessly query databases. In particular, here we focus on improving the structure grounding ability of text-to-SQL models, which is critical in real-world use cases. We evaluate our model with the widely used Spider benchmark and several more realistic datasets. Experimental results show that our method brings significant improvement over existing baselines, especially on more realistic settings.

Computing cost. We use 4 V100 GPUs for pre-training, and 1 V100 GPU for finetuning the model for text-to-SQL on Spider and WikiSQL. One advantage of our method is its efficiency. In our experiments, we pretrain the model for only 5 epochs, which can finish within 4 hours. For comparison, the largest TaBERT model (Yin et al., 2020) takes 6 days to train for 10 epochs on 128 Tesla V100 GPUs using mixed precision training.

Acknowledgements

We thank Bo Pang, Tao Yu for their help with the official Spider evaluation. We also thank anonymous reviewers for their constructive feedback.

References

Ion Androutsopoulos, G. Ritchie, and P. Thanisch. 1995. Natural language interfaces to databases - an introduction. Nat. Lang. Eng., 1:29–81.

Ben Bogin, Matt Gardner, and Jonathan Berant. 2019. Global reasoning over database structures for text-to-SQL parsing. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 3659–3664, Hong Kong, China. Association for Computational Linguistics.

Wenhu Chen, Jianshu Chen, Yu Su, Zhiyu Chen, and William Yang Wang. 2020a. Logical natural language generation from open-domain tables. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 7929–7942, Online. Association for Computational Linguistics.

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Yang Wang. 2020b. HybridQA: A dataset of multi-hop question answering over tabular and textual data. In Findings of the Association for Computational Linguistics: EMNLP 2020, pages 1026–1036, Online. Association for Computational Linguistics.

DongHyun Choi, Myeong Cheol Shin, EungGyun Kim, and Dong Ryeol Shin. 2020. Ryansql: Recursively applying sketch-based slot fillings for complex text-to-sql in cross-domain databases. arXiv preprint arXiv:2004.03125.

Deborah A. Dahl, Madeleine Bates, Michael Brown, William Fisher, Kate Hunicke-Smith, David Pallett, Christine Pao, Alexander Rudnicky, and Elizabeth Shriberg. 1994. Expanding the scope of the ATIS task: The ATIS-3 corpus. In Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, Volume 1 (Long and Short Papers), pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Catherine Finegan-Dollak, Jonathan K. Kummerfeld, Li Zhang, Karthik Ramanathan, Sesh Sadasivam, Rui Zhang, and Dragomir Radev. 2018. Improving text-to-SQL evaluation methodology. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 351–360, Melbourne, Australia. Association for Computational Linguistics.

Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. 2019. Towards complex text-to-SQL in cross-domain database with intermediate representation. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 4524–4535, Florence, Italy. Association for Computational Linguistics.

Pengcheng He, Yi Mao, Kaushik Chakrabarti, and Weizhu Chen. 2019. X-sql: reinforce schema representation with context. arXiv preprint arXiv:1908.08113.

Charles T. Hemphill, John J. Godfrey, and George R. Doddington. 1990. The ATIS spoken language systems pilot corpus. In Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27,1990.

Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. TaPas: Weakly supervised table parsing via pre-training. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 4320–4333, Online. Association for Computational Linguistics.

Wonseok Hwang, Jinyeong Yim, Seunghyun Park, and Minjoon Seo. 2019. A comprehensive exploration on wikisql with table-aware word contextualization. arXiv preprint arXiv:1902.01069.

Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. Learning a neural semantic parser from user feedback. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 963–973, Vancouver, Canada. Association for Computational Linguistics.

Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings.

Oliver Lehmberg, Dominique Ritze, Robert Meusel, and Christian Bizer. 2016. A large public corpus of web tables containing time and context metadata. In Proceedings of the 25th International Conference Companion on World Wide Web, pages 75–76.

Wenqiang Lei, Weixin Wang, Zhixin Ma, Tian Gan, Wei Lu, Min-Yen Kan, and Tat-Seng Chua. 2020. Re-examining the role of schema linking in text-to-SQL. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 6943–6954, Online. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 7871–7880, Online. Association for Computational Linguistics.

F. Li and H. V. Jagadish. 2014a. Nalir: an interactive natural language interface for querying relational databases. In SIGMOD Conference.

Fei Li and HV Jagadish. 2014b. Constructing an interactive natural language interface for relational databases. Proceedings of the VLDB Endowment, 8(1):73–84.

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

Qin Lyu, Kaushik Chakrabarti, Shobhit Hathi, Souvik Kundu, Jianwen Zhang, and Zheng Chen. 2020. Hybrid ranking network for text-to-sql. Technical Report MSR-TR-2020-7, Microsoft Dynamics 365 AI.

Ankur Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. 2020. ToTTo: A controlled table-to-text generation dataset. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 1173–1186, Online. Association for Computational Linguistics.

Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 1470–1480, Beijing, China. Association for Computational Linguistics.

Tianze Shi, Chen Zhao, Jordan Boyd-Graber, Hal Daumé III, and Lillian Lee. 2020. On the potential of lexico-logical alignments for semantic parsing to SQL queries. In Findings of the Association for Computational Linguistics: EMNLP 2020, pages 1849–1864, Online. Association for Computational Linguistics.

Alane Suhr, Ming-Wei Chang, Peter Shaw, and Kenton Lee. 2020. Exploring unexplored generalization challenges for cross-database semantic parsing. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 8372–8388, Online. Association for Computational Linguistics.

Lappoon R. Tang and Raymond J. Mooney. 2000. Automated construction of database interfaces: Intergrating statistical and relational learning for semantic parsing. In 2000 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora, pages 133–141, Hong Kong, China. Association for Computational Linguistics.

Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020. 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, pages 7567–7578, Online. Association for Computational Linguistics.

Chenglong Wang, Kedar Tatwawadi, Marc Brockschmidt, Po-Sen Huang, Yi Mao, Oleksandr Polozov, and Rishabh Singh. 2018. Robust text-to-sql generation with execution-guided decoding. arXiv preprint arXiv:1807.03100.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 38–45, Online. Association for Computational Linguistics.

Navid Yaghmazadeh, Yuepeng Wang, Isil Dillig, and Thomas Dillig. 2017. Sqlizer: query synthesis from natural language. Proc. ACM Program. Lang., 1(OOPSLA):63:1–63:26.

Ziyu Yao, Yiqi Tang, Wen-tau Yih, Huan Sun, and Yu Su. 2020. An imitation game for learning semantic parsers from user interaction. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pages 6883–6902, Online. Association for Computational Linguistics.

Semih Yavuz, Izzeddin Gur, Yu Su, and Xifeng Yan. 2018. What it takes to achieve 100% condition accuracy on WikiSQL. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pages 1702–1711, Brussels, Belgium. Association for Computational Linguistics.

Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. TaBERT: Pretraining for joint understanding of textual and tabular data. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 8413–8426, Online. Association for Computational Linguistics.

Tao Yu, Zifan Li, Zilin Zhang, Rui Zhang, and Dragomir Radev. 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, Volume 2 (Short Papers), pages 588–594, New Orleans, Louisiana. Association for Computational Linguistics.

Tao Yu, Chien-Sheng Wu, Xi Victoria Lin, Bailin Wang, Yi Chern Tan, Xinyi Yang, Dragomir Radev, Richard Socher, and Caiming Xiong. 2020. Grappa: Grammar-augmented pre-training for table semantic parsing. arXiv preprint arXiv:2009.13845.

Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018b. 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, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.

John M Zelle and Raymond J Mooney. 1996. Learning to parse database queries using inductive logic programming. In Proceedings of the national conference on artificial intelligence, pages 1050–1055.

Rui Zhang, Tao Yu, Heyang Er, Sungrok Shim, Eric Xue, Xi Victoria Lin, Tianze Shi, Caiming Xiong, Richard Socher, and Dragomir Radev. 2019. Editing-based SQL query generation for cross-domain context-dependent questions. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 5338–5349, Hong Kong, China. Association for Computational Linguistics.

Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. CoRR, abs/1709.00103.

Models	Easy	Medium	Hard	Extra Hard	All
# Examples	248	446	174	166	1034
RAT-SQL *w/o* value linking					
w. BERT$_{\text{LARGE}}$	82.9	72.7	65.7	46.6	69.8
w. STRUG (Human Assisted)	84.9	76.2	67.0	55.0	73.3
w. STRUG (Automatic)	87.8	75.6	69.5	55.0	74.2
RAT-SQL					
w. BERT$_{\text{LARGE}}$	84.1	74.9	67.4	52.8	72.3
w. STRUG (Human Assisted)	87.1	77.7	70.9	57.0	75.5
w. STRUG (Automatic)	88.7	77.4	69.2	53.6	74.9

Table 7: Execution accuracy on Spider dev set with different hardness levels.

Models	SELECT	WHERE	GROUP BY	ORDER BY
RAT-SQL *w/o* value linking				
w. BERT$_{\text{LARGE}}$	89.2	71.7	78.7	81.5
w. STRUG (Human Assisted)	91.2	74.8	79.0	84.0
w. STRUG (Automatic)	90.9	75.6	77.5	84.0
RAT-SQL				
w. BERT$_{\text{LARGE}}$	89.4	79.2	78.5	81.3
w. STRUG (Human Assisted)	91.3	80.8	80.6	85.7
w. STRUG (Automatic)	91.2	80.1	78.6	84.5

Table 8: F1 scores of Component Matching on Spider dev set.

A Implementation Details

A.1 Filtering on the Suhr et al. (2020) Datasets

We use the filtering scripts[10] released by the authors of Suhr et al. (2020). More specifically, they remove examples that fall into the following categories: (1) a numeric or text value in the query is not copiable from the utterance (except for the numbers 0 and 1, which are often not copied from the input), (2) the result of the query is a empty table, or a query for count returns [1], (3) the query requires selecting more than one final column.

B More Resutls

B.1 Detailed Results on Spider and Spider-Realistic

We show more detailed results on the Spider dev set and Spider-Realistic in Table 7, Table 8 and Table 9. From Table 7 we can see that STRUG brings significant improvements in all difficulty levels, and is not biased towards certain subset. Since STRUG mostly improves the structure grounding ability of the model, from Table 8 and Table 9, we can see that STRUG mainly increase the accuracy for WHERE and ORDER BY clauses, especially when database content is not available to the model. On the Spider-Realistic set, as the model cannot rely on simple string matching for structure grounding,

---

10 https://github.com/google-research/language/tree/master/language/xsp

Models	SELECT	WHERE	GROUP BY	ORDER BY
RAT-SQL *w/o* value linking				
w. BERT$_{\text{LARGE}}$	86.2	55.6	65.9	64.3
w. STRUG (Human Assisted)	88.9	61.9	70.4	64.1
w. STRUG (Automatic)	90.1	64.5	73.0	67.4
RAT-SQL				
w. BERT$_{\text{LARGE}}$	86.9	74.2	59.6	61.9
w. STRUG (Human Assisted)	89.0	76.8	69.9	63.5
w. STRUG (Automatic)	89.2	76.4	64.7	64.9

Table 9: F1 scores of Component Matching on Spider-Realistic set.

Models	ACC$_{\text{S-COL}}$	ACC$_{\text{S-AGG}}$	ACC$_{\text{W-COL}}$	ACC$_{\text{W-VAL}}$
SQLova (5%)				
w. BERT$_{\text{LARGE}}$	95.2	88.4	89.6	88.3
w. TaBERT	95.4	88.4	90.8	88.0
w. STRUG (Human Assisted)	95.5	88.9	92.6	91.5
w. STRUG (Automatic)	95.8	88.9	92.3	91.7

Table 10: Subtask performance on WikiSQL. S-COL, S-AGG, W-COL and W-VAL stands for tasks of predicting SELECT column, aggregation operator, WHERE columns and WHERE values, respectively.

we notice greater improvement using STRUG, especially for GROUP BY clauses.

B.2 Detailed Results on WikiSQL

We show subtask performance for WikiSQL in Table 10, Fig. 7 and Fig. 4. Again, we can see that STRUG mainly improves WHERE column and WHERE value accuracy. From Fig. 6 we can see that with only 1% of training data, model with STRUG already has over 0.87 WHERE column accuracy and nearly 0.85 WHERE value accuracy.
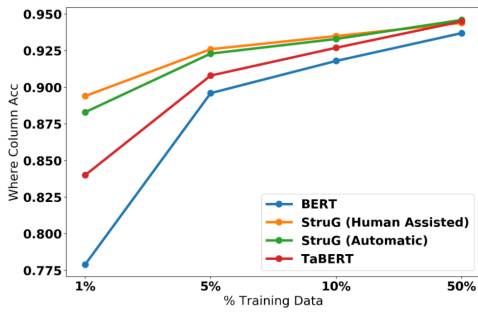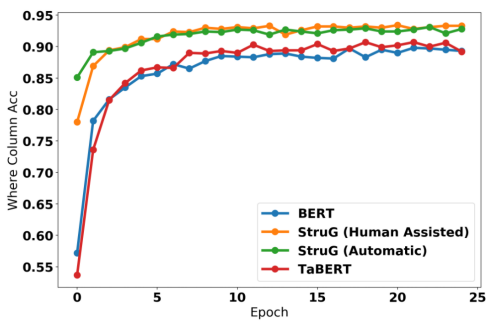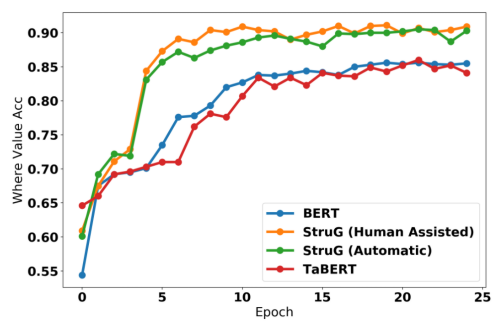
(a) Where Column Accuracy
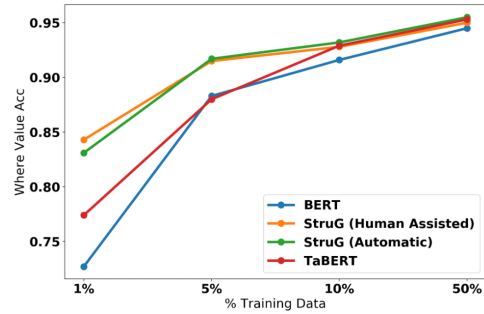
(b) Where Value Accuracy

Figure 6: Model performance on the test set with different fractions of training data.

(a) Where Column Accuracy

(b) Where Value Accuracy

Figure 7: Model performance on the dev set during training with 5% of training data.