

Semantic Parsing with Syntax- and Table-Aware SQL Generation

Yibo Sun^{§*}, Duyu Tang[‡], Nan Duan[‡], Jianshu Ji[‡], Guihong Cao[‡],
Xiaocheng Feng[§], Bing Qin[§], Ting Liu[§], Ming Zhou[‡]

[§]Harbin Institute of Technology, Harbin, China

[‡]Microsoft Research Asia, Beijing, China

[‡]Microsoft AI and Research, Redmond WA, USA

{ybsun, xcfeng, qinb, tliu}@ir.hit.edu.cn

{dutang, nanduan, jianshuj, gucao, mingzhou}@microsoft.com

Abstract

We present a generative model to map natural language questions into SQL queries. Existing neural network based approaches typically generate a SQL query word-by-word, however, a large portion of the generated results is incorrect or not executable due to the mismatch between question words and table contents. Our approach addresses this problem by considering the structure of table and the syntax of SQL language. The quality of the generated SQL query is significantly improved through (1) learning to replicate content from column names, cells or SQL keywords; and (2) improving the generation of WHERE clause by leveraging the column-cell relation. Experiments are conducted on WikiSQL, a recently released dataset with the largest question-SQL pairs. Our approach significantly improves the state-of-the-art execution accuracy from 69.0% to 74.4%.

1 Introduction

We focus on semantic parsing that maps natural language utterances to executable programs (Zelle and Mooney, 1996; Wong and Mooney, 2007; Zettlemoyer and Collins, 2007; Kwiatkowski et al., 2011; Pasupat and Liang, 2015; Iyer et al., 2017; Iyyer et al., 2017). In this work, we regard SQL as the programming language, which could be executed on a table or a relational database to obtain an outcome. Datasets are the main driver of progress for statistical approaches in semantic parsing (Liang, 2016). Recently, Zhong

et al. (2017) release WikiSQL, the largest hand-annotated semantic parsing dataset which is an order of magnitude larger than other datasets in terms of both the number of logical forms and the number of tables. Pointer network (Vinyals et al., 2015) based approach is developed, which generates a SQL query word-by-word through replicating from a word sequence consisting of question words, column names and SQL keywords. However, a large portion of generated results is incorrect or not executable due to the mismatch between question words and column names (or cells). This also reflects the real scenario where users do not always use exactly the same column name or cell content to express the question.

To address the aforementioned problem, we present a generative semantic parser that considers the structure of table and the syntax of SQL language. The approach is partly inspired by the success of structure/grammar driven neural network approaches in semantic parsing (Xiao et al., 2016; Krishnamurthy et al., 2017). Our approach is based on pointer networks, which encodes the question into continuous vectors, and synthesizes the SQL query with three channels. The model learns when to generate a column name, a cell or a SQL keyword. We further incorporate column-cell relation to mitigate the ill-formed outcomes.

We conduct experiments on WikiSQL. Results show that our approach outperforms existing systems, improving state-of-the-art execution accuracy to 74.4% and logical form accuracy to 60.7%. An extensive analysis reveals the advantages and limitations of our approach.

2 Task Formulation and Dataset

As shown in Figure 1, we focus on sequence-to-SQL generation in this work. Formally, the task takes a question q and a table t consisting of n col-

* Work is done during internship at Microsoft Research Asia.

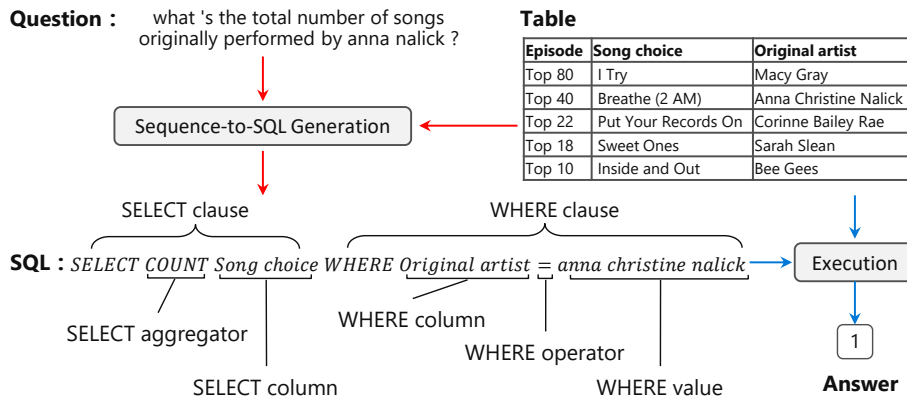


Figure 1: An brief illustration of the task. The focus of this work is sequence-to-SQL generation.

umn names and $n \times m$ cells as the input, and outputs a SQL query y . We do not consider the join operation over multiple relational tables, which we leave in the future work.

We use WikiSQL (Zhong et al., 2017), the largest hand-annotated semantic parsing dataset to date which consists of 87,726 questions and SQL queries distributed across 26,375 tables from Wikipedia.

3 Related Work

Semantic Parsing. Semantic parsing aims to map natural language utterances to programs (e.g., logical forms), which will be executed to obtain the answer (denotation) (Zettlemoyer and Collins, 2005; Liang et al., 2011; Berant et al., 2013; Poon, 2013; Krishnamurthy and Kollar, 2013; Pasupat and Liang, 2016; Sun et al., 2016; Jia and Liang, 2016; Kočiský et al., 2016; Lin et al., 2017). Existing studies differ from (1) the form of the knowledge base, e.g. facts from Freebase, a table (or relational database), an image (Suhr et al., 2017; Johnson et al., 2017; Hu et al., 2017; Goldman et al., 2017) or a world state (Long et al., 2016); (2) the programming language, e.g. first-order logic, lambda calculus, lambda DCS, SQL, parameterized neural programmer (Yin et al., 2015; Neelakantan et al., 2016), or coupled distributed and symbolic executors (Mou et al., 2017); (3) the supervision used for learning the semantic parser, e.g. question-denotation pairs, binary correct/incorrect feedback (Artzi and Zettlemoyer, 2013), or richer supervision of question-logical form pairs (Dong and Lapata, 2016). In this work, we study semantic parsing over tables, which is critical for users to access relational databases with natural language, and could serve users' in-

formation need for structured data on the web. We use SQL as the programming language, which has a broad acceptance to programmers.

Natural Language Interface for Databases.

Our work relates to the area of accessing database with natural language interface (Dahl et al., 1994; Brad et al., 2017). Popescu et al. (2003) use a parser to parse the question, and then use lexicon matching between question and the table column names/cells. Giordani and Moschitti (2012) parse the question with dependency parser, compose candidate SQL queries with heuristic rules, and use kernel based SVM ranker to rank the results. Li and Jagadish (2014) translate natural language utterances into SQL queries based on dependency parsing results, and interact with users to ensure the correctness of the interpretation process. Yaghmazadeh et al. (2017) build a semantic parser on the top of SEMPRES (Pasupat and Liang, 2015) to get a SQL sketch, which only has the SQL shape and will be subsequently completed based on the table content. Iyer et al. (2017) maps utterances to SQL queries through sequence-to-sequence learning. User feedbacks are incorporated to reduce the number of queries to be labeled. Zhong et al. (2017) develop an augmented pointer network, which is further improved with reinforcement learning for SQL sequence prediction. Xu et al. (2017) adopts a sequence-to-set model to predict WHERE columns, and uses an attentional model to predict the slots in where clause.

Different from (Iyer et al., 2017; Zhong et al., 2017), our approach leverages SQL syntax and table structure. Compared to (Popescu et al., 2003; Giordani and Moschitti, 2012; Yaghmazadeh et al., 2017), our approach is end-to-end learning and independent of a syntactic parser or manu-

ally designed templates. We are aware of existing studies that combine reinforcement learning and maximum likelihood estimation (MLE) (Guu et al., 2017; Mou et al., 2017; Liang et al., 2017). However, the focus of this work is the design of the neural architecture, despite we also implement an RL strategy (refer to §4.4).

Structure/Grammar Guided Neural Decoder

Our approach could be viewed as an extension of the sequence-to-sequence learning (Sutskever et al., 2014; Bahdanau et al., 2015) with a tailored neural decoder driven by the characteristic of the target language (Yin and Neubig, 2017; Rabinovich et al., 2017). Methods with similar intuitions have been developed for language modeling (Dyer et al., 2016), neural machine translation (Wu et al., 2017) and lambda calculus based semantic parsing (Dong and Lapata, 2016; Krishnamurthy et al., 2017). The difference is that our model is developed for sequence-to-SQL generation, in which table structure and SQL syntax are considered.

4 Methodology

We first describe the background on pointer networks, and then present our approach that considers the table structure and the SQL syntax.

4.1 Background: Pointer Networks

Pointer networks is originally introduced by (Vinyals et al., 2015), which takes a sequence of elements as the input and outputs a sequence of discrete tokens corresponding to positions in the input sequence. The approach has been successfully applied in reading comprehension (Kadlec et al., 2016) for pointing to the positions of answer span from the document, and also in sequence-to-sequence based machine translation (Gulcehre et al., 2016) and text summarization (Gu et al., 2016) for replicating rare words from the source sequence to the target sequence.

The approach of Zhong et al. (2017) is based on pointer networks. The encoder is a recurrent neural network (RNN) with gated recurrent unit (GRU) (Cho et al., 2014), whose input is the concatenation of question words, words from column names and SQL keywords. The decoder is another GRU based RNN, which works in a sequential way and generates a word at each time step. The generation of a word is actually selectively replicating a word from the input sequence, the prob-

ability distribution of which is calculated with an attention mechanism (Bahdanau et al., 2015). The probability of generating the i -th word x_i in the input sequence at the t -th time step is calculated as Equation 1, where h_t^{dec} is the decoder hidden state at the t -th time step, h_i^{enc} is the encoder hidden state of the word x_i , W_a is the model parameter.

$$p(y_t = x_i | y_{<t}, x) \propto \exp(W_a [h_t^{dec}; h_i^{enc}]) \quad (1)$$

It is worth to note that if a column name consists of multiple words (such as “original artist” in Figure 1), these words are separated in the input sequence. The approach has no guarantee that a multi-word column name could be successively generated, which would affect the executability of the generated SQL query.

4.2 STAMP: Syntax- and Table- Aware seMantic Parser

Figure 2 illustrates an overview of the proposed model, which is abbreviated as STAMP. Different from Zhong et al. (2017), the word is not the basic unit to be generated in STAMP. As is shown, there are three “channels” in STAMP, among which the column channel predicts a column name, the value channel predicts a table cell and the SQL channel predicts a SQL keyword. Accordingly, the probability of generating a target token is formulated in Equation 2, where z_t stands for the channel selected by the switching gate, $p_z(\cdot)$ is the probability to choose a channel, and $p_w(\cdot)$ is similar to Equation 1 which is a probability distribution over the tokens from one of the three channels.

$$p(y_t | y_{<t}, x) = \sum_{z_t} p_w(y_t | z_t, y_{<t}, x) p_z(z_t | y_{<t}, x) \quad (2)$$

One advantage of this architecture is that it inherently addresses the problem of generating partial column name/cell because an entire column name/cell is the basic unit to be generated. Another advantage is that the column-cell relation and question-cell connection can be naturally integrated in the model, which will be described below.

Specifically, our encoder takes a question as the input. Bidirectional RNN with GRU unit is applied to the question, and the concatenation of both ends is used as the initial state of the decoder. Another bidirectional RNN is used to compute the representation of a column name (or a cell), in case that each unit contains multiple words (Dong

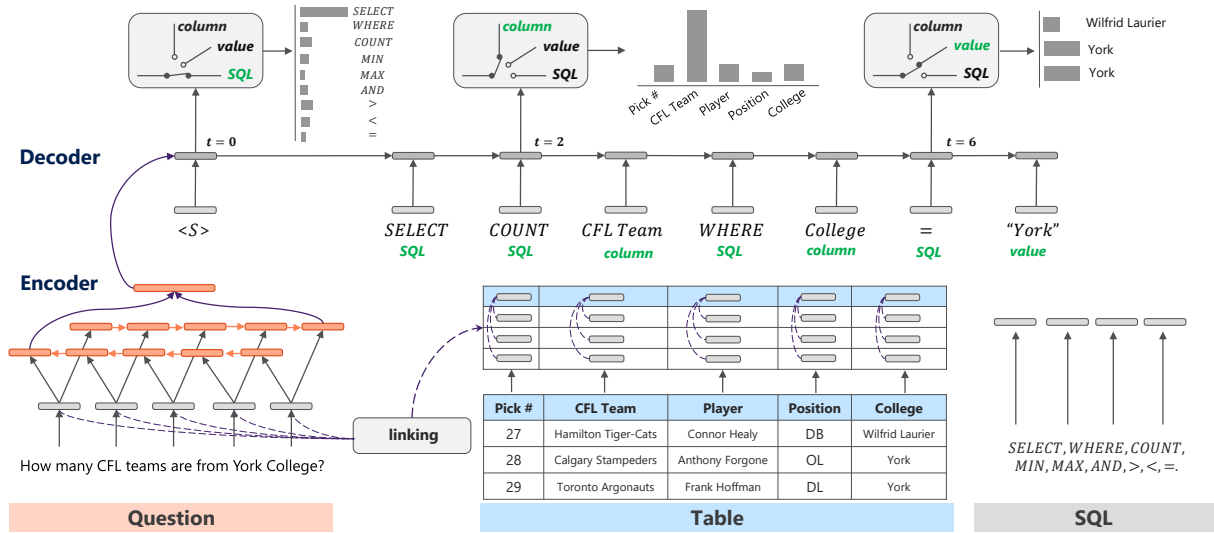


Figure 2: An illustration of the proposed approach. At each time step, a switching gate selects a channel to predict a column name (maybe composed of multiple words), a cell or a SQL keyword. The words in green below the SQL tokens stand for the results of the switching gate at each time step.

et al., 2015). Essentially, each channel is an attentional neural network. For cell and SQL channels, the input of the attention module only contains the decoder hidden state and the representation of the token to be calculated as follows,

$$p_w^{sql}(i) \propto \exp(W_{sql}[h_t^{dec}; e_i^{sql}]) \quad (3)$$

where e_i^{sql} stands for the representation of the i -th SQL keyword. As suggested by (Zhong et al., 2017), we also concatenate the question representation into the input of the column channel in order to improve the accuracy of the SELECT column. We implement the switching gate with a feed-forward neural network, in which the output is a *softmax* function and the input is the decoder hidden state h_t^{dec} .

4.3 Improved with Column-Cell Relation

We further improve the STAMP model by considering the column-cell relation, which is important for predicting the WHERE clause.

On one hand, the column-cell relation could improve the prediction of SELECT column. We observe that a cell or a part of it typically appears at the question acting as the WHERE value, such as "anna nalick" for "anna christine nalick"). However, a column name might be represented with a totally different utterance, which is a "semantic gap". Supposing the question is "How many schools did player number 3 play at?" and the SQL query is "Select count School

Club Team where No. = 3". We could see that the column names "School Club Team" and "No." are different from their corresponding utterances "schools", "number" in natural language question. Thus, table cells could be regarded as the pivot that connects the question and column names (the "linking" component in Figure 2). For instance, taking the question from Figure 2, the word "York" would help to predict the column name as "College" rather than "Player". There might be different possible ways to implement this intuition. We use cell information to enhance the column name representation in this work. The vector of a column name is further concatenated with a question-aware cell vector, which is weighted averaged over the cell vectors belonging to the same column. The probability distribution in the column channel is calculated as Equation 4. We use the number of cell words occurring in the question to measure the importance of a cell, which is further normalized through a *softmax* function to yield the final weight $\alpha_j^{cell} \in [0, 1]$. An alternative measurement is to use an additional attention model whose input contains the question vector and the cell vector. We favor to the intuitive and efficient way in this work.

$$p_w^{col}(i) \propto \exp(W_{col}[h_t^{dec}; h_i^{col}; \sum_{j \in col_i} \alpha_j^{cell} h_j^{cell}]) \quad (4)$$

On the other hand, the column-cell relation could improve the prediction of the WHERE val-

ue. To yield an executable SQL, the WHERE value should be a cell that belongs to the same WHERE column¹. Taking Figure 2 as an example, it should be avoided to predict a where clause like “*Player = York*” because the cell “*York*” does not belong to the column name “*Player*”. To achieve this, we incorporate a global variable to memorize the last predicted column name. When the switching gate selects the value channel, the cell distribution is only calculated over the cells belonging to the last predicted column name. Furthermore, we incorporate an additional probability distribution over cells based on the aforementioned word co-occurrence between the question and cells, and weighted average two cell distributions, which is calculated as follows.

$$p_w^{cell}(j) = \lambda \hat{p}_w^{cell}(j) + (1 - \lambda) \alpha_j^{cell} \quad (5)$$

where $\hat{p}_w^{cell}(j)$ is the standard probability distribution obtained from the attentional neural network, and λ is a hyper parameter which is tuned on the dev set.

4.4 Improved with Policy Gradient

The model described so far could be conventionally learned via cross-entropy loss over question-SQL pairs. However, different SQL queries might be executed to yield the same result, and possible SQL queries of different variations could not be exhaustively covered in the training dataset. Two possible ways to handle this are (1) shuffling the WHERE clause to generate more SQL queries, and (2) using reinforcement learning (RL) which regards the correctness of the executed output as the goodness (reward) of the generated SQL query. We follow Zhong et al. (2017) and adopt a policy gradient based approach. We use a baseline strategy (Zaremba and Sutskever, 2015) to decrease the learning variance. The expected reward (Williams, 1992) for an instance is calculated as $\mathbb{E}(y_g) = \sum_{j=1}^k \log p(y_j) R(y_j, y^g)$, where y^g is the ground truth SQL query, y_j is a generated SQL query, $p(y_j)$ is the probability of y_j being generated by our model, and k is the number of sampled SQL queries. $R(y_j, y^g)$ is the same reward function defined by Zhong et al. (2017), which is +1 if y_j is executed to yield the correct answer; -1 if

¹This constraint is suitable in this work as we do not consider the nested query in the where clause, such as “*where College = select College from table*”, which is also the case not included in the WikiSQL dataset. We leave generating nested SQL query in the future work.

y_j is a valid SQL query and is executed to get an incorrect answer; and -2 if y_j is not a valid SQL query. In this way, model parameters could be updated with policy gradient over question-answer pairs.

4.5 Training and Inference

As the WikiSQL data contains rich supervision of question-SQL pairs, we use them to train model parameters. The model has two cross-entropy loss functions, as given below. One is for the switching gate classifier (p_z) and another is for the attentional probability distribution of a channel (p_w).

$$l = - \sum_t \log p_z(z_t | y_{<t}, x) - \sum_t \log p_w(y_t | z_t, y_{<t}, x) \quad (6)$$

Our parameter setting strictly follows Zhong et al. (2017). We represent each word using word embedding² (Pennington et al., 2014) and the mean of the sub-word embeddings of all the n-grams in the word (Hashimoto et al., 2016)³. The dimension of the concatenated word embedding is 400. We clamp the embedding values to avoid over-fitting. We set the dimension of encoder and decoder hidden state as 200. During training, we randomize model parameters from a uniform distribution with fan-in and fan-out, set batch size as 64, set the learning rate of SGD as 0.5, and update the model with stochastic gradient descent. Greedy search is used in the inference process. We use the model trained from question-SQL pairs as initialization and use RL strategy to fine-tune the model. SQL queries used for training RL are sampled based on the probability distribution of the model learned from question-SQL pairs. We tune the best model on the dev set and do inference on the test set for only once. This protocol is used in model comparison as well as in ablations.

5 Experiment

We conduct experiments on the WikiSQL dataset⁴, which includes 61, 297/9, 145/17, 284 examples in the training/dev/test sets. Each instance consists of a question, a table, a SQL query and a result. Following Zhong et al. (2017), we use two

²<http://nlp.stanford.edu/data/glove.840B.300d.zip>

³http://www.logos.t.u-tokyo.ac.jp/~hassy/publications/arxiv2016jmt/jmt_pre-trained_embeddings.tar.gz

⁴<https://github.com/salesforce/WikiSQL>

Methods	Dev		Test	
	Acc _{lf}	Acc _{ex}	Acc _{lf}	Acc _{ex}
Attentional Seq2Seq	23.3%	37.0%	23.4%	35.9%
Aug.PntNet (Zhong et al., 2017)	44.1%	53.8%	43.3%	53.3%
Aug.PntNet (re-implemented by us)	51.5%	58.9%	52.1%	59.2%
Seq2SQL (no RL) (Zhong et al., 2017)	48.2%	58.1%	47.4%	57.1%
Seq2SQL (Zhong et al., 2017)	49.5%	60.8%	48.3%	59.4%
SQLNet (Xu et al., 2017)	–	69.8%	–	68.0%
Guo and Gao (2018)	–	71.1%	–	69.0%
STAMP (w/o cell)	58.6%	67.8%	58.0%	67.4%
STAMP (w/o column-cell relation)	59.3%	71.8%	58.4%	70.6%
STAMP	61.5%	74.8%	60.7%	74.4%
STAMP+RL	61.7%	75.1%	61.0%	74.6%

Table 1: Performances of different approaches on the WikiSQL dataset. Two evaluation metrics are logical form accuracy (Acc_{lf}) and execution accuracy (Acc_{ex}). Our model is abbreviated as (STAMP).

evaluation metrics. One metric is logical form accuracy (Acc_{lf}), which measures the percentage of the generated SQL queries that have exact string match with the ground truth SQL queries. Since different SQL queries might obtain the same result, another metric is execution accuracy (Acc_{ex}), which measures the percentage of the generated SQL queries that obtain the correct answer.

5.1 Model Comparisons

After released, WikiSQL dataset has attracted a lot of attentions from both industry and research communities. Zhong et al. (2017) develop following methods, including (1) Aug.PntNet which is an end-to-end learning pointer network approach; (2) Seq2SQL (no RL), in which two separate classifiers are trained for SELECT aggregator and SELECT column, separately; and (3) Seq2SQL, in which reinforcement learning is further used for model training. Results of attentional sequence-to-sequence learning baseline (Attentional Seq2Seq) are also reported in (Zhong et al., 2017). Xu et al. (2017) develop SQLNet, which predicts SELECT clause and WHERE clause separately. Sequence-to-set neural architecture and column attention are adopted to predict the WHERE clause. Similarly, Guo and Gao (2018) develop tailored modules to handle three components of SQL queries, respectively. A parallel work from (Yu et al., 2018) obtains higher execution accuracy (82.6%) on WikiSQL, however, its model is slot-filling based which is designed specifically for the “select-aggregator-where” type and utilizes external knowledge base (such as Freebase) to tag the

question words. We believe this mechanism could improve our model as well, we leave this as a potential future work.

Our model is abbreviated as (STAMP), which is short for Syntax- and Table- Aware seMantic Parser. The STAMP model in Table 1 stands for the model we describe in §4.2 plus §4.3. STAMP+RL is the model that is fine-tuned with the reinforcement learning strategy as described in §4.4. We implement a simplified version of our approach (w/o cell), in which WHERE values come from the question. Thus, this setting differs from Aug.PntNet in the generation of WHERE column. We also study the influence of the relation-cell relation (w/o column-cell relation) through removing the enhanced column vector, which is calculated by weighted averaging cell vectors.

From Table 1, we can see that STAMP performs better than existing systems on WikiSQL. Incorporating RL strategy does not significantly improve the performance. Our simplified model, STAMP (w/o cell), achieves better accuracy than Aug.PntNet, which further reveals the effects of the column channel. Results also demonstrate the effects of incorporating the column-cell relation, removing which leads to about 4% performance drop in terms of Acc_{ex}.

5.2 Model Analysis: Fine-Grained Accuracy

We analyze the STAMP model from different perspectives in this part.

Firstly, since SQL queries in WikiSQL consists of SELECT column, SELECT aggregator, and WHERE clause, we report the results with regard

Methods	Dev			Test		
	Acc _{sel}	Acc _{agg}	Acc _{where}	Acc _{sel}	Acc _{agg}	Acc _{where}
Aug.PntNet (reimplemented by us)	80.9%	89.3%	62.1%	81.3%	89.7%	62.1%
Seq2SQL (Zhong et al., 2017)	89.6%	90.0%	62.1%	88.9%	90.1%	60.2%
SQLNet (Xu et al., 2017)	91.5%	90.1%	74.1%	90.9%	90.3%	71.9%
Guo and Gao (2018)	92.5%	90.1%	74.7%	91.9%	90.3%	72.8%
STAMP (w/o cell)	89.9%	89.2%	72.1%	89.2%	89.3%	71.2%
STAMP (w/o column-cell relation)	89.3%	89.2%	73.2%	88.8%	89.2%	71.8%
STAMP	89.4%	89.5%	77.1%	88.9%	89.7%	76.0%
STAMP+RL	89.6%	89.7%	77.3%	90.0%	89.9%	76.3%

Table 2: Fine-grained accuracies on the WikiSQL dev and test sets. Accuracy (Acc_f) is evaluated on SELECT column (Acc_{sel}), SELECT aggregator (Acc_{agg}), and WHERE clause (Acc_{where}), respectively.

to more fine-grained evaluation metrics over these aspects. Results are given in Table 2, in which the numbers of Seq2SQL and SQLNet are reported in (Xu et al., 2017). We can see that the main improvement of STAMP comes from the WHERE clause, which is also the key challenge of the WikiSQL dataset. This is consistent with our primary intuition on improving the prediction of WHERE column and WHERE value. The accuracies of STAMP on SELECT column and SELECT aggregator are not as high as SQLNet. The main reason is that these two approaches train the SELECT clause separately while STAMP learns all these components in a unified paradigm.

5.3 Model Analysis: Difficulty Analysis

We study the performance of STAMP on different portions of the test set according to the difficulties of examples. We compare between Aug.PntNet (re-implemented by us) and STAMP. In this work, the difficulty of an example is reflected by the number of WHERE columns.

Method	#where	Dev	Test
Aug.PntNet	= 1	63.4%	63.8%
	= 2	51.0%	51.8%
	≥ 3	38.5%	38.1%
STAMP	= 1	80.9%	80.2%
	= 2	65.1%	65.4%
	≥ 3	44.1%	48.2%

Table 3: Execution accuracy (Acc_{ex}) on different groups of WikiSQL dev and test sets.

From Table 3, we can see that STAMP outperforms Aug.PntNet in all these groups. The accuracy decreases with the increase of the number of WHERE conditions.

5.4 Model Analysis: Executable Analysis

We study the percentage of executable SQL queries in the generated results. As shown in Table 4, STAMP significantly outperforms Aug.PntNet. Almost all the results of STAMP are executable. This is because STAMP avoids generating incomplete column names or cells, and guarantees the correlation between WHERE conditions and WHERE values in the table.

	Dev	Test
Aug.PntNet	77.9%	78.7%
STAMP	99.9%	99.9%

Table 4: Percentage of the executable SQL queries on WikiSQL dev and test sets.

5.5 Model Analysis: Case Study

We give a case study to illustrate the generated results by STAMP, with a comparison to Aug.PntNet. Results are given in Figure 3. In the first example, Aug.PntNet generates incomplete column name (“*style*”), which is addressed in STAMP through replicating an entire column name. In the second example, the WHERE value (“*brazilian jiu-jitsu*”) does not belong to the generated WHERE column (“*Masters*”) in Aug.PntNet. This problem is avoided in STAMP through incorporating the table content.

5.6 Error Analysis

We conduct error analysis on the dev set of WikiSQL to show the limitation of the STAMP model and where is the room for making further improvements. We analyze the 2,302 examples which are executed to wrong answers by the STAMP model, and find that 33.6% of them have wrong SE-

Episode #	Country	City	Martial Art/Style	Masters	Original Airdate
1.1	China	Dengfeng	Kung Fu (Wushu ; Sanda)	Shi De Yang, Shi De Cheng	28-Dec-07
1.2	Philippines	Manila	Kali	Leo T. Gaje Jr. Cristino Vasquez	4-Jan-08
1.3	Japan	Tokyo	Kyokushin Karate	Yuzo Goda, Isamu Fukuda	11-Jan-08
1.4	Mexico	Mexico City	Boxing	Ignacio "Nacho" Beristáin Tiburcio Garcia	18-Jan-08
1.5	Indonesia	Bandung	Pencak Silat	Rita Suwanda Dadang Gunawan	25-Jan-08
1.7	South Korea	Seoul	Hapkido	Kim Nam Je, Bae Sung Book Ju Soong Weo	8-Feb-08
1.8	Brazil	Rio de Janeiro	Brazilian Jiu-Jitsu	Breno Sivak, Renato Barreto Royler Gracie	15-Feb-08
1.9	Israel	Netanya	Krav Maga	Ran Nakash Avivit Oftek Cohen	22-Feb-08

Question #1: how many masters fought using a boxing style ?
Aug.PntNet: select count masters from table where style = boxing
STAMP: select count masters from table where martial art/style = boxing

Question #2: when did the episode featuring a master using brazilian jiu-jitsu air ?
Aug.PntNet: select original airdate from table where masters = brazilian jiu-jitsu
STAMP: select original airdate from table where martial art/style = brazilian jiu-jitsu

Figure 3: Case study on the dev set between Aug.PntNet and STAMP. These two questions are based on the same table. Each question is followed by the generated SQL queries from the two approaches.

LECT columns, 15.7% of them have a different number of conditions in the WHERE clause, and 53.7% of them have a different WHERE column set compared to the ground truth. Afterwards, we analyze a portion of randomly sampled dissatisfied examples. Consistent with the qualitative results, most problems come from column prediction, including both SELECT clause and WHERE clause. Even though the overall accuracy of the SELECT column prediction is about 90% and we also use cell information to enhance the column representation, this semantic gap is still the main bottleneck. Extracting and incorporating various expressions for a table column (i.e. relation in a relational database) might be a potential way to mitigate this problem. Compared to column prediction, the quality of cell prediction is much better because cell content typically (partially) appears in the question.

5.7 Transfers to WikiTableQuestions

WikiTableQuestions (Pasupat and Liang, 2015) is a widely used dataset for semantic parsing. To further test the performance of our approach, we conduct an additional transfer learning experiment. Firstly, we directly apply the STAMP model trained on WikiSQL to WikiTableQuestions, which is an unsupervised learning setting for the WikiTableQuestions dataset. Results show that the test accuracy of STAMP in this setting is 14.5%, which has a big gap between best systems on WikiTableQuestions, where Zhang et al. (2017)

and Krishnamurthy et al. (2017) yield 43.3% and 43.7%, respectively. Furthermore, we apply the learnt STAMP model to generate SQL queries on natural language questions from WikiTableQuestions, and regard the generated SQL queries which could be executed to correct answers as additional pseudo question-SQL pairs. In this way, the STAMP model learnt from a combination of WikiSQL and pseudo question-SQL pairs could achieve 21.0% on the test set. We find that this big gap is caused by the difference between the two datasets. Among 8 types of questions in WikiTableQuestions, half of them including {"Union", "Intersection", "Reverse", "Arithmetic"} are not covered in the WikiSQL dataset. It is an interesting direction to leverage algorithms developed from two datasets to improve one another.

5.8 Discussion

Compared to slot-filling based models that restrict target SQL queries to fixed forms of "select-aggregator-where", our model is less tailored. We believe that it is easy to expand our model to generate nested SQL queries or JOIN clauses, which could also be easily trained with back-propagation if enough training instances of these SQL types are available. For example, we could incorporate a hierarchical "value" channel to handle nest queries. Let us suppose our decoder works horizontally that next generated token is at the right hand of the current token. Inspired by chunk-based decoder for neural machine translation (Ishiwatari et al.,

2017), we could increase the depth of the “value” channel to generate tokens of a nested WHERE value along the vertical axis. During inference, an addition gating function might be necessary to determine whether to generate a nested query, followed by the generation of WHERE value. An intuitive way that extends our model to handle JOIN clauses is to add the 4th channel, which predicts a table from a collection of tables. Therefore, the decoder should learn to select one of the four channels at each time step. Accordingly, we need to add “from” as a new SQL keyword in order to generate SQL queries including “from xxxTable”.

In terms of the syntax of SQL, the grammar we used in this work could be regarded as shallow syntax, such as three channels and column-cell relation. We do not use deep syntax, such as the sketch of SQL language utilized in some slot-filling models, because incorporating them would make the model clumpy. Instead, we let the model to learn the sequential and compositional relations of SQL queries automatically from data. Empirical results show that our model learns these patterns well.

6 Conclusion and Future Work

In this work, we develop STAMP, a Syntax- and Table- Aware seMantic Parser that automatically maps natural language questions to SQL queries, which could be executed on web table or relational dataset to get the answer. STAMP has three channels, and it learns to switch to which channel at each time step. STAMP considers cell information and the relation between cell and column name in the generation process. Experiments are conducted on the WikiSQL dataset. Results show that STAMP achieves the new state-of-the-art performance on WikiSQL. We conduct extensive experiment analysis to show advantages and limitations of our approach, and where is the room for others to make further improvements.

SQL language has more complicated queries than the cases included in the WikiSQL dataset, including (1) querying over multiple relational databases, (2) nested SQL query as condition value, (3) more operations such as “group by” and “order by”, etc. In this work, the STAMP model is not designed for the first and second cases, but it could be easily adapted to the third case through incorporating additional SQL keywords and of course the learning of which requires dataset of the

same type. In the future, we plan to improve the accuracy of the column prediction component. We also plan to build a large-scale dataset that considers more sophisticated SQL queries. We also plan to extend the approach to low-resource scenarios (Feng et al., 2018).

Acknowledgments

We thank Yaming Sun for her great help. We also would like to thank three anonymous reviewers for their valuable comments and suggestions. This research was partly supported by National Natural Science Foundation of China(No. 61632011 and No.61772156, and No.61472107).

References

- Yoav Artzi and Luke Zettlemoyer. 2013. Weakly supervised learning of semantic parsers for mapping instructions to actions. *Transactions of the Association for Computational Linguistics* 1:49–62.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. *Proceeding of ICLR*.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *EMNLP*. 5, page 6.
- Florin Brad, Radu Cristian Alexandru Iacob, Ionel Alexandru Hosu, and Traian Rebedea. 2017. Dataset for a neural natural language interface for databases (nlibdb). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Asian Federation of Natural Language Processing, Taipei, Taiwan, pages 906–914. <http://www.aclweb.org/anthology/I17-1091>.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, pages 1724–1734. <http://www.aclweb.org/anthology/D14-1179>.
- 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 *Proceedings of the workshop on Human Language Technology*. Association for Computational Linguistics, pages 43–48.

- Li Dong and Mirella Lapata. 2016. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 33–43. <http://www.aclweb.org/anthology/P16-1004>.
- Li Dong, Furu Wei, Hong Sun, Ming Zhou, and Ke Xu. 2015. A hybrid neural model for type classification of entity mentions. In *IJCAI*. pages 1243–1249.
- Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. 2016. Recurrent neural network grammars. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, San Diego, California, pages 199–209. <http://www.aclweb.org/anthology/N16-1024>.
- Xiaocheng Feng, Xiachong Feng, Bing Qin, Zhangyin Feng, and Ting Liu. 2018. Improving low resource named entity recognition using cross-lingual knowledge transfer. In *IJCAI*.
- Alessandra Giordani and Alessandro Moschitti. 2012. Translating questions to sql queries with generative parsers discriminatively reranked. In *COLING (Posters)*. pages 401–410.
- Omer Goldman, Veronica Latcinnik, Udi Naveh, Amir Globerson, and Jonathan Berant. 2017. Weakly-supervised semantic parsing with abstract examples. *CoRR* abs/1711.05240. <http://arxiv.org/abs/1711.05240>.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 1631–1640. <http://www.aclweb.org/anthology/P16-1154>.
- Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the unknown words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 140–149. <http://www.aclweb.org/anthology/P16-1014>.
- Tong Guo and Huilin Gao. 2018. Bidirectional attention for SQL generation. *CoRR* abs/1801.00076. <http://arxiv.org/abs/1801.00076>.
- Kelvin Guu, Panupong Pasupat, Evan Liu, and Percy Liang. 2017. From language to programs: Bridging reinforcement learning and maximum marginal likelihood. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. pages 1051–1062.
- Kazuma Hashimoto, Caiming Xiong, Yoshimasa Tsuruoka, and Richard Socher. 2016. A joint many-task model: Growing a neural network for multiple nlp tasks. *arXiv preprint arXiv:1611.01587*.
- Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Kate Saenko. 2017. Learning to reason: End-to-end module networks for visual question answering. *International Conference on Computer Vision (ICCV)*.
- Shonosuke Ishiwatari, Jingtao Yao, Shujie Liu, Mu Li, Ming Zhou, Naoki Yoshinaga, Masaru Kitsuregawa, and Weijia Jia. 2017. Chunk-based decoder for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vancouver, Canada, pages 1901–1912. <http://aclweb.org/anthology/P17-1174>.
- 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)*. Association for Computational Linguistics, Vancouver, Canada, pages 963–973. <http://aclweb.org/anthology/P17-1089>.
- Mohit Iyyer, Wen-tau Yih, and Ming-Wei Chang. 2017. Search-based neural structured learning for sequential question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vancouver, Canada, pages 1821–1831. <http://aclweb.org/anthology/P17-1167>.
- Robin Jia and Percy Liang. 2016. Data recombination for neural semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 12–22. <http://www.aclweb.org/anthology/P16-1002>.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Judy Hoffman, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Inferring and executing programs for visual reasoning. *International Conference on Computer Vision (ICCV)*.
- Rudolf Kadlec, Martin Schmid, Ondřej Bajgar, and Jan Kleindienst. 2016. Text understanding with the attention sum reader network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 908–918. <http://www.aclweb.org/anthology/P16-1086>.
- Tomáš Kočiský, Gábor Melis, Edward Grefenstette, Chris Dyer, Wang Ling, Phil Blunsom, and Karl Moritz Hermann. 2016. Semantic parsing with

- semi-supervised sequential autoencoders. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, pages 1078–1087. <https://aclweb.org/anthology/D16-1116>.
- Jayant Krishnamurthy, Pradeep Dasigi, and Matt Gardner. 2017. Neural semantic parsing with type constraints for semi-structured tables. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, pages 1517–1527. <https://www.aclweb.org/anthology/D17-1160>.
- Jayant Krishnamurthy and Thomas Kollar. 2013. Jointly learning to parse and perceive: Connecting natural language to the physical world. *Transactions of the Association for Computational Linguistics* 1:193–206.
- Tom Kwiatkowski, Luke Zettlemoyer, Sharon Goldwater, and Mark Steedman. 2011. Lexical generalization in ccg grammar induction for semantic parsing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pages 1512–1523.
- Fei Li and HV Jagadish. 2014. Constructing an interactive natural language interface for relational databases. *Proceedings of the VLDB Endowment* 8(1):73–84.
- Chen Liang, Jonathan Berant, Quoc Le, Kenneth D. Forbus, and Ni Lao. 2017. Neural symbolic machines: Learning semantic parsers on freebase with weak supervision. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vancouver, Canada, pages 23–33. <http://aclweb.org/anthology/P17-1003>.
- Percy Liang. 2016. Learning executable semantic parsers for natural language understanding. *Communications of the ACM* 59(9):68–76.
- Percy Liang, Michael I Jordan, and Dan Klein. 2011. Learning dependency-based compositional semantics. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*. pages 590–599.
- Xi Victoria Lin, Chenglong Wang, Deric Pang, Kevin Vu, Luke Zettlemoyer, and Michael D. Ernst. 2017. Program synthesis from natural language using recurrent neural networks. Technical Report UW-CSE-17-03-01, University of Washington Department of Computer Science and Engineering, Seattle, WA, USA.
- Reginald Long, Panupong Pasupat, and Percy Liang. 2016. Simpler context-dependent logical forms via model projections. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 1456–1465. <http://www.aclweb.org/anthology/P16-1138>.
- Lili Mou, Zhengdong Lu, Hang Li, and Zhi Jin. 2017. Coupling distributed and symbolic execution for natural language queries. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*. pages 2518–2526. <http://proceedings.mlr.press/v70/mou17a.html>.
- Arvind Neelakantan, Quoc V Le, Martin Abadi, Andrew McCallum, and Dario Amodei. 2016. Learning a natural language interface with neural programmer. *arXiv preprint arXiv:1611.08945*.
- 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)*. Association for Computational Linguistics, Beijing, China, pages 1470–1480. <http://www.aclweb.org/anthology/P15-1142>.
- Panupong Pasupat and Percy Liang. 2016. Inferring logical forms from denotations. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. pages 23–32.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Doha, Qatar, pages 1532–1543. <http://www.aclweb.org/anthology/D14-1162>.
- Hoifung Poon. 2013. Grounded unsupervised semantic parsing. In *ACL (1)*. pages 933–943.
- Ana-Maria Popescu, Oren Etzioni, and Henry Kautz. 2003. Towards a theory of natural language interfaces to databases. In *Proceedings of the 8th international conference on Intelligent user interfaces*. ACM, pages 149–157.
- Maxim Rabinovich, Mitchell Stern, and Dan Klein. 2017. Abstract syntax networks for code generation and semantic parsing. *arXiv preprint arXiv:1704.07535*.
- Alane Suhr, Mike Lewis, James Yeh, and Yoav Artzi. 2017. A corpus of natural language for visual reasoning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Vancouver, Canada, pages 217–223. <http://aclweb.org/anthology/P17-2034>.

- Huan Sun, Hao Ma, Xiaodong He, Wen-tau Yih, Yu Su, and Xifeng Yan. 2016. Table cell search for question answering. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, pages 771–782.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. pages 3104–3112.
- Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*. pages 2692–2700.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning* 8(3-4):229–256.
- Yuk Wah Wong and Raymond J Mooney. 2007. Learning synchronous grammars for semantic parsing with lambda calculus. In *Annual Meeting-Association for computational Linguistics*. 1, page 960.
- Shuangzhi Wu, Dongdong Zhang, Nan Yang, Mu Li, and Ming Zhou. 2017. Sequence-to-dependency neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vancouver, Canada, pages 698–707. <http://aclweb.org/anthology/P17-1065>.
- Chunyang Xiao, Marc Dymetman, and Claire Gardent. 2016. Sequence-based structured prediction for semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 1341–1350. <http://www.aclweb.org/anthology/P16-1127>.
- Xiaojun Xu, Chang Liu, and Dawn Song. 2017. Sqlnet: Generating structured queries from natural language without reinforcement learning. *arXiv preprint arXiv:1711.04436*.
- Navid Yaghmazadeh, Yuepeng Wang, Isil Dillig, and Thomas Dillig. 2017. Type-and content-driven synthesis of sql queries from natural language. *arXiv preprint arXiv:1702.01168*.
- Pengcheng Yin, Zhengdong Lu, Hang Li, and Ben Kao. 2015. Neural enquirer: Learning to query tables with natural language. *arXiv preprint arXiv:1512.00965*.
- Pengcheng Yin and Graham Neubig. 2017. A syntactic neural model for general-purpose code generation. *arXiv preprint arXiv:1704.01696*.
- Tao Yu, Zifan Li, Zilin Zhang, Rui Zhang, and Dragomir Radev. 2018. Typesql: Knowledge-based type-aware neural text-to-sql generation. *arXiv preprint arXiv:1804.09769*.
- Wojciech Zaremba and Ilya Sutskever. 2015. Reinforcement learning neural turing machines. *arXiv preprint arXiv:1505.00521* 419.
- 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.
- Luke S. Zettlemoyer and Michael Collins. 2005. Learning to map sentences to logical form: Structured classification with probabilistic categorial grammars. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence*. pages 658–666.
- Luke S Zettlemoyer and Michael Collins. 2007. Online learning of relaxed ccg grammars for parsing to logical form. In *EMNLP-CoNLL*. pages 678–687.
- Yuchen Zhang, Panupong Pasupat, and Percy Liang. 2017. Macro grammars and holistic triggering for efficient semantic parsing. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, pages 1225–1234. <https://www.aclweb.org/anthology/D17-1126>.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.