

RAT-SQL: RELATION-AWARE SCHEMA ENCODING AND LINKING FOR TEXT-TO-SQL PARSERS

Anonymous authors

Paper under double-blind review

ABSTRACT

When translating natural language questions into SQL queries to answer questions from a database, contemporary semantic parsing models struggle to generalize to unseen database schemas. The generalization challenge lies in (a) encoding the database relations in an accessible way for the semantic parser, and (b) modeling alignment between database columns and their mentions in a given query. We present a unified framework, based on the relation-aware self-attention mechanism, to address schema encoding, schema linking, and feature representation within a text-to-SQL encoder. On the challenging Spider dataset this framework boosts the exact match accuracy to 53.7%, compared to 47.4% for the state-of-the-art model unaugmented with BERT embeddings. In addition, we observe qualitative improvements in the model’s understanding of schema linking and alignment.

1 INTRODUCTION

The ability to effectively query databases with natural language has the potential to unlock the power of large datasets to the vast majority of users who are not proficient in query languages. As such, a large body of research has focused on the task of translating natural language questions into queries that existing database software can execute.

The release of large annotated datasets containing questions and the corresponding database SQL queries has catalyzed progress in the field, by enabling the training of supervised learning models for the task. In contrast to prior semantic parsing datasets (Finegan-Dollak et al., 2018), new tasks such as WikiSQL (Zhong et al., 2017) and Spider (Yu et al., 2018b) pose the real-life challenge of *generalization to unseen database schemas*. Every query is conditioned on a multi-table database schema, and the databases do not overlap between the train and test sets.

Schema generalization is challenging for three interconnected reasons. First, any text-to-SQL semantic parsing model must encode a given schema into column and table representations suitable for decoding a SQL query that might involve any of the given columns or tables. Second, these representations should encode all the information about the schema, including its column types, foreign key relations, and primary keys used for database joins. Finally, the model must recognize natural language used to refer to database columns and tables, which might differ from the referential language seen in training. The latter challenge is known as *schema linking* – aligning column/table references in the question to the corresponding schema columns/tables.




Figure 1: A challenging text-to-SQL task from the Spider dataset.

While the question of *schema encoding* has been studied in recent literature (Bogin et al., 2019b), *schema linking* has been relatively less explored. Consider the example in Figure 1. It illustrates the challenge of ambiguity in linking: while “*model*” in the question refers to `car_names.model` rather than `model_list.model`, “*cars*” actually refers to both `cars_data` and `car_names` (but not `car_makers`) for the purpose of table joining. To resolve the column/table references properly, the semantic parser must take into account both the known schema relations (e.g. foreign keys) and the question context.

Prior work (Bogin et al., 2019b) addressed the schema representation problem by encoding the directed graph of foreign key relations among the columns with a graph neural network. While effective, this approach has two important shortcomings. First, it does not contextualize schema encoding with the question, thus making it difficult for the model to reason about schema linking after both the column representations and question word representations have been built. Second, it limits information propagation during schema encoding to predefined relations in the schema such as foreign keys. The advent of self-attentional mechanisms in natural language processing (Vaswani et al., 2017) shows that global reasoning is crucial to building effective representations of relational structures. However, we would like any global reasoning to also take into account the aforementioned predefined schema relations.

In this work, we present a unified framework, called RAT-SQL,¹ for encoding relational structure in the database schema and a given question. It uses *relation-aware self-attention* to combine global reasoning over the schema entities and question words with structured reasoning over predefined schema relations. We then apply RAT-SQL to the problems of schema encoding and schema linking. As a result, we obtain 53.7% exact match accuracy on the Spider test set. At the time of writing, this result is the state of the art among models unaugmented with pretrained BERT embeddings. In addition, we experimentally demonstrate that RAT-SQL enables the model to build more accurate internal representations of the question’s true alignment with schema columns and tables.

2 RELATED WORK

Semantic parsing of natural language to SQL queries recently surged in popularity thanks to the creation of two new multi-table datasets with the challenge of schema generalization – WikiSQL (Zhong et al., 2017) and Spider (Yu et al., 2018b). Schema encoding is not as challenging in WikiSQL as in Spider thanks to the lack of multi-table relations. Schema linking is relevant for both tasks but also more challenging in Spider due to the richer natural language expressiveness and less restricted SQL grammar observed in it. Indeed, the state of the art semantic parser on WikiSQL (He et al., 2019) achieves a test set accuracy of 91.8%, significantly higher than the state of the art on Spider.

The recent state-of-the-art models evaluated on Spider use various attentional architectures for question/schema encoding and AST-based structural architectures for query decoding. IRNet (Guo et al., 2019) encodes the question and schema separately with LSTM and self-attention respectively, augmenting them with custom type vectors for schema linking. They further use the AST-based decoder of Yin and Neubig (2017) to decode a query in an intermediate representation (IR) that exhibits higher-level abstraction structure than SQL. Bogin et al. (2019b) encode the schema with a graph neural network and a similar grammar-based decoder. Both approaches highlight the importance of schema encoding and schema linking, but design separate feature engineering techniques to augment *word vectors* (as opposed to *relations between words and columns*) to resolve it. In contrast, the relational framework of RAT-SQL provides a unified way to encode arbitrary relational information among the inputs.

Concurrently with this work, Bogin et al. (2019a) published Global-GNN, a different approach to schema linking for Spider which applies global reasoning between question words and schema columns/tables. Global reasoning is implemented by *gating* the graph neural network that computes the representation of schema elements using question token representations. This conceptually differs from RAT-SQL in two important ways: (a) question word representations influence the schema representations but not vice versa, and (b) like in other GNN-based encoding approaches, message propagation is limited to the schema-induced edges such as foreign key relations. In contrast, our

¹Relation-Aware Transformer.




Figure 2: An illustration of an example schema as a graph. We do not depict all edges and label types of Table 1 to reduce clutter.

relation-aware transformer mechanism allows encoding arbitrary relations between question words and schema elements explicitly, and these representations are computed jointly using self-attention.

We use the same formulation of relation-aware self-attention as Shaw et al. (2018). However, that work only applied it to sequences of words in the context of machine translation, and as such, their set of relation types only encoded the relative distance between two words. We extend their work and show that relation-aware self-attention can effectively encode more complex relationships that exist within an unordered sets of elements (in this case, columns and tables within a database schema as well as relations between the schema and the question). To the best of our knowledge, this is the first application of relation-aware self-attention to joint representation learning with both predefined and softly induced relations in the input structure.

3 RAT-SQL

We now describe the RAT-SQL framework and its application to the problems of schema encoding and linking. First, we formally define the text-to-SQL semantic parsing problem and its components. Then, we introduce the relation-aware self-attention mechanism, our framework for jointly encoding relational structure between the question and the schema. Finally, we present our implementation of schema linking in the RAT-SQL framework.

3.1 PROBLEM DEFINITION

Given a natural language question Q and a schema $S = \langle \mathcal{C}, \mathcal{T} \rangle$ for a relational database, our goal is to generate the corresponding SQL P . Here the question $Q = q_1 \dots q_{|Q|}$ is a sequence of words, and the schema consists of columns $\mathcal{C} = \{c_1, \dots, c_{|\mathcal{C}|}\}$ and tables $\mathcal{T} = \{t_1, \dots, t_{|\mathcal{T}|}\}$. Each column name c_i contains words $c_{i,1}, \dots, c_{i,|c_i|}$ and each table name t_i contains words $t_{i,1}, \dots, t_{i,|t_i|}$. The desired program P is represented as an *abstract syntax tree* T in the context-free grammar of SQL.

Some columns in the schema are *primary keys*, used for uniquely indexing the corresponding table, and some are *foreign keys*, used to reference a primary key column in a different table. As described in Section 1, we would like to softly bias our schema encoding mechanism toward these predefined relations. In addition, each column has a *type* τ such as `number` or `text`.

Schema linking aims at finding the alignment between question words and mentioned columns or tables. It’s a crucial step for a parser to generate the right columns and tables in SQL. We model the latent alignment explicitly using an alignment matrix (Section 3.6), which is softly biased towards some string-match based relations, as inspired by Guo et al. (2019).

3.2 ENCODING THE SCHEMA AS A GRAPH

To support reasoning about relationships between schema elements in the encoder, we begin by representing the database schema using a directed graph \mathcal{G} , where each node and edge has a label. We represent each table and column in the schema as a node in this graph, labeled with the words in the name; for columns, we prepend the type of the column to the label. For each pair of nodes x and y in the graph, Table 1 describes when there exists an edge from x to y and the label it should have. Figure 2 illustrates an example graph (although not all edges and labels are shown).

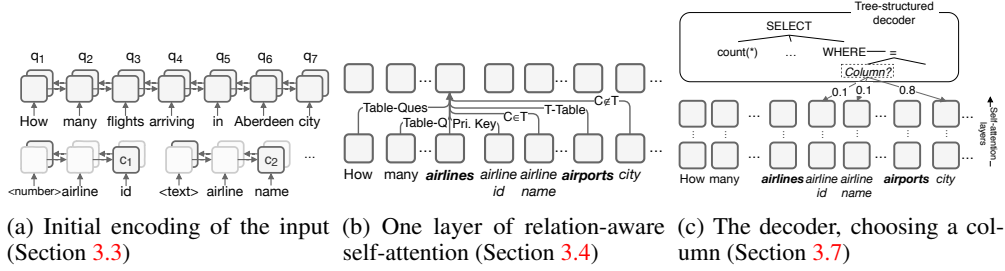


Figure 3: Overview of the stages of our approach.

3.3 INITIAL ENCODING OF THE INPUT

We now obtain an initial representation for each of the nodes in the graph, as well as for the words in the input question. For the graph nodes, we use a bidirectional LSTM (BiLSTM) over the words contained in the label. We concatenate the output of the initial and final time steps of this LSTM to form the embedding for the node. For the question, we also use a bidirectional LSTM over the words:

$$\begin{aligned}
 (\mathbf{c}_{i,0}^{\text{fwd}}, \mathbf{c}_{i,0}^{\text{rev}}, \dots, (\mathbf{c}_{i,|c_i|}^{\text{fwd}}, \mathbf{c}_{i,|c_i|}^{\text{rev}})) &= \text{BiLSTM}_{\text{Column}}(c_i^{\text{type}}, c_{i,1}, \dots, c_{i,|c_i|}); & \mathbf{c}_i^{\text{init}} &= \text{Concat}(\mathbf{c}_{i,|c_i|}^{\text{fwd}}, \mathbf{c}_{i,0}^{\text{rev}}) \\
 (\mathbf{t}_{i,1}^{\text{fwd}}, \mathbf{t}_{i,1}^{\text{rev}}, \dots, (\mathbf{t}_{i,|t_i|}^{\text{fwd}}, \mathbf{t}_{i,|t_i|}^{\text{rev}})) &= \text{BiLSTM}_{\text{Table}}(t_{i,1}, \dots, t_{i,|t_i|}); & \mathbf{t}_i^{\text{init}} &= \text{Concat}(\mathbf{t}_{i,|t_i|}^{\text{fwd}}, \mathbf{t}_{i,1}^{\text{rev}}) \\
 (\mathbf{q}_1^{\text{fwd}}, \mathbf{q}_1^{\text{rev}}, \dots, (\mathbf{q}_{|Q|}^{\text{fwd}}, \mathbf{q}_{|Q|}^{\text{rev}})) &= \text{BiLSTM}_{\text{Question}}(q_1, \dots, q_{|Q|}); & \mathbf{q}_i^{\text{init}} &= \text{Concat}(\mathbf{q}_i^{\text{fwd}}, \mathbf{q}_i^{\text{rev}})
 \end{aligned}$$

where each of the BiLSTM functions first lookup word embeddings for each of the input tokens. The LSTMs do not share any parameters.

3.4 RELATION-AWARE TRANSFORMER

At this point, we have representations $\mathbf{c}_i^{\text{init}}$, $\mathbf{t}_i^{\text{init}}$, and $\mathbf{q}_i^{\text{init}}$. Similar to encoders used in some previous papers, these initial representations are independent of each other (uninfluenced by which other columns or tables are present). Now, we would like to imbue these representations with the information in the schema graph. We use a form of self-attention (Vaswani et al., 2017) that is relation-aware (Shaw et al., 2018) to achieve this goal.

In one step of relation-aware self-attention, we begin with an input \mathbf{x} of n elements (where $x_i \in \mathbb{R}^{d_x}$) and transform each x_i into $y_i \in \mathbb{R}^{d_x}$. We follow the formulation described in Shaw et al. (2018):

$$\begin{aligned}
 e_{ij}^{(h)} &= \frac{x_i W_Q^{(h)} (x_j W_K^{(h)} + \mathbf{r}_{ij}^K)^T}{\sqrt{d_z/H}}; & \alpha_{ij}^{(h)} &= \frac{\exp(e_{ij}^{(h)})}{\sum_{l=1}^n \exp(e_{il}^{(h)})} \\
 z_i^{(h)} &= \sum_{j=1}^n \alpha_{ij}^{(h)} (x_j W_V^{(h)} + \mathbf{r}_{ij}^V); & z_i &= \text{Concat}(z_i^{(1)}, \dots, z_i^{(H)}) \\
 \tilde{y}_i &= \text{LayerNorm}(x_i + z_i); & y_i &= \text{LayerNorm}(\tilde{y}_i + \text{FC}(\text{ReLU}(\text{FC}(\tilde{y}_i))))
 \end{aligned}$$

where FC is a fully-connected layer, $1 \leq h \leq H$, and $W_Q^{(h)}, W_K^{(h)}, W_V^{(h)} \in \mathbb{R}^{d_x \times (d_x/H)}$. The r_{ij} terms encode the relationship between the two elements x_i and x_j in the input. We explain how we obtain r_{ij} in the next part.

Application Within Our Encoder At the start, we construct the input \mathbf{x} of $|c| + |t| + |q|$ elements using $\mathbf{c}_i^{\text{init}}$, $\mathbf{t}_i^{\text{init}}$, and $\mathbf{q}_i^{\text{init}}$:

$$\mathbf{x} = (\mathbf{c}_1^{\text{init}}, \dots, \mathbf{c}_{|c|}^{\text{init}}, \mathbf{t}_1^{\text{init}}, \dots, \mathbf{t}_{|t|}^{\text{init}}, \mathbf{q}_1^{\text{init}}, \dots, \mathbf{q}_{|q|}^{\text{init}}).$$

We then apply a stack of N relation-aware self-attention layers, where N is a hyperparameter. The weights of the encoder layers are not tied; each layer has its own set of weights. After processing through the stack of N encoder layers, we obtain

$$(\mathbf{c}_1^{\text{final}}, \dots, \mathbf{c}_{|c|}^{\text{final}}, \mathbf{t}_1^{\text{final}}, \dots, \mathbf{t}_{|t|}^{\text{final}}, \mathbf{q}_1^{\text{final}}, \dots, \mathbf{q}_{|q|}^{\text{final}}) = \mathbf{y}.$$

Table 1: Description of edge types present in the directed graph created to represent the schema. An edge exists from source node $x \in \mathcal{S}$ to target node $y \in \mathcal{S}$ if the pair fulfills one of the descriptions listed in the table, with the corresponding label. Otherwise, no edge exists from x to y .

Type of x	Type of y	Edge label	Description
Column	Column	SAME-TABLE	x and y belong to the same table.
		FOREIGN-KEY-COL-F	x is a foreign key for y .
		FOREIGN-KEY-COL-R	y is a foreign key for x .
Column	Table	PRIMARY-KEY-F	x is the primary key of y .
		BELONGS-TO-F	x is a column of y (but not the primary key).
Table	Column	PRIMARY-KEY-R	y is the primary key of x .
		BELONGS-TO-R	y is a column of x (but not the primary key).
Table	Table	FOREIGN-KEY-TAB-F	Table x has a foreign key column in y .
		FOREIGN-KEY-TAB-R	Same as above, but x and y are reversed.
		FOREIGN-KEY-TAB-B	x and y have foreign keys in both directions.

We use c_i^{final} , t_i^{final} , and q_i^{final} in our decoder.

We define a discrete set of possible relation types, and map each type to an embedding to obtain r_{ij}^V and r_{ij}^K . We need a value of r_{ij} for every pair of elements in x . In the subsequent sections, we describe the set of relation types we used.

3.5 SCHEMA ENCODING

If x_i and x_j both correspond to nodes in \mathcal{G} (i.e. each is either a column or table) with an edge from x_i to x_j , then we use the label on that edge (possibilities listed in Table 1) for r_{ij} . However, this is not sufficient to obtain r_{ij} for every pair of i and j . The graph \mathcal{G} has no nodes corresponding to the question words, not every pair of schema nodes has an edge between them, and there is no self-edges (for when $i = j$). As such, we add more types beyond what is defined in Table 1:

- If $i = j$, then COLUMN-IDENTITY or TABLE-IDENTITY.
- $x_i \in \text{question}$, $x_j \in \text{question}$: QUESTION-DIST- d , where $d = \text{clip}(j - i, D)$; $\text{clip}(a, D) = \max(-D, \min(D, a))$. We use $D = 2$.
- $x_i \in \text{question}$, $x_j \in \text{column} \cup \text{table}$; or $x_i \in \text{column} \cup \text{table}$, $x_j \in \text{question}$: see Section 3.6.
- Otherwise, one of COLUMN-COLUMN, COLUMN-TABLE, TABLE-COLUMN, or TABLE-TABLE.

3.6 SCHEMA LINKING

To aid the model with aligning column/table references in the question to the corresponding schema columns/tables, we furthermore define relation types which indicate when parts of the question textually match the names of the columns and tables. Specifically, for all n-grams of length 1 to 5 in the question, we determine (1) whether it exactly matches the name of a column/table (*exact match*); or (2) whether the n-gram is a subsequence of the name of a column/table (*partial match*).²

Therefore, for the case where $x_i \in \text{question}$, $x_j \in \text{column} \cup \text{table}$; or $x_i \in \text{column} \cup \text{table}$, $x_j \in \text{question}$, we set r_{ij} to QUESTION-COLUMN-M, QUESTION-TABLE-M, COLUMN-QUESTION-M or TABLE-QUESTION-M depending on the type of x_i and x_j . M is one of EXACTMATCH, PARTIALMATCH, or NOMATCH. In the end, we add $2 + 5 + (4 \times 3) + 4$ types (one term per bullet in Section 3.5) beyond the 10 in Table 1, for a total of 33 types.

Memory-Schema Alignment Matrix Our intuition suggests that the columns and tables which occur in the SQL P will generally have a corresponding reference in the natural language question (for example, “cars” and “cylinders” in Figure 1). To capture this intuition in the model, we apply

²This procedure matches that of Guo et al. (2019), but we use the matching information differently in RAT.

relation-aware attention as a pointer mechanism between every memory element in y and all the columns/tables to compute explicit *alignment matrices* $L^{\text{col}} \in \mathbb{R}^{|y| \times |C|}$ and $L^{\text{tab}} \in \mathbb{R}^{|y| \times |\mathcal{T}|}$:

$$\begin{aligned}\tilde{L}_{i,j}^{\text{col}} &= \frac{y_i W_Q^{\text{col}} (\mathbf{c}_j^{\text{final}} W_K^{\text{col}} + \mathbf{r}_{ij}^K)^T}{\sqrt{d_x}}; & L_{i,j}^{\text{col}} &= \frac{\exp(\tilde{L}_{i,j}^{\text{col}})}{\sum_{k=1}^{|C|} \exp(\tilde{L}_{i,k}^{\text{col}})} \\ \tilde{L}_{i,j}^{\text{tab}} &= \frac{y_i W_Q^{\text{tab}} (\mathbf{t}_j^{\text{final}} W_K^{\text{tab}} + \mathbf{r}_{ij}^K)^T}{\sqrt{d_x}}; & L_{i,j}^{\text{tab}} &= \frac{\exp(\tilde{L}_{i,j}^{\text{tab}})}{\sum_{k=1}^{|\mathcal{T}|} \exp(\tilde{L}_{i,k}^{\text{tab}})}\end{aligned}$$

The memory-schema alignment matrix is expected to resemble the real discrete alignments, therefore should respect certain constraints like sparsity. For example, the question word “model” in Figure 1 should be aligned with `car_names.model` rather than `model_list.model` or `model_list.model_id`. To further bias the soft alignment towards the real discrete structures, we add an auxiliary loss to encourage sparsity of the alignment matrix. Specifically, for a column/table that is mentioned in the SQL query, we treat the model’s current belief of the best alignment as the ground truth. Then we use a cross-entropy loss, referred as *alignment loss*, to strengthen the model’s belief:

$$\text{align_loss} = -\frac{1}{|Rel(C)|} \sum_{j \in Rel(C)} \log \max_i L_{i,j}^{\text{col}} - \frac{1}{|Rel(\mathcal{T})|} \sum_{j \in Rel(\mathcal{T})} \log \max_i L_{i,j}^{\text{tab}}$$

where $Rel(C)$ and $Rel(\mathcal{T})$ denote the set of *relevant* columns and tables that appear in the SQL P .

3.7 DECODER

Once we have obtained an encoding of the input, we used the decoder from Yin and Neubig (2017) to generate the SQL P . The decoder generates P as an abstract syntax tree in depth-first traversal order, by using an LSTM to output a sequence of *decoder actions* that (i) expand the last generated node in the tree according to the grammar, called `APPLYRULE`; or when necessary to complete the last node, (ii) chooses a column or table from the schema, called `SELECTCOLUMN` and `SELECTTABLE`. Formally, we have the following:

$$\Pr(P | y) = \prod_t \Pr(a_t | a_{<t}, y)$$

where y is the final encoding of the question and schema from the previous section, and $a_{<t}$ are all previous actions. We update the LSTM’s state in the following way: $\mathbf{m}_t, \mathbf{h}_t = f_{\text{LSTM}}([\mathbf{a}_{t-1} \| \mathbf{z}_t \| \mathbf{h}_{p_t} \| \mathbf{a}_{p_t} \| \mathbf{n}_{f_t}], \mathbf{m}_{t-1}, \mathbf{h}_{t-1})$ where \mathbf{m}_t is the LSTM cell state, \mathbf{h}_t is the LSTM output at step t , \mathbf{a}_{t-1} is the embedding of the previous action, p_t is the step corresponding to expanding the parent AST node of the current node, and \mathbf{n}_{f_t} is the embedding of the current node type. We obtain \mathbf{z}_t using multi-head attention (with 8 heads) on \mathbf{h}_{t-1} over y .

For `APPLYRULE`[R], we compute $\Pr(a_t = \text{APPLYRULE}[R] | a_{<t}, y) = \text{softmax}_R(g(\mathbf{h}_t))$ where $g(\cdot)$ is a 2-layer MLP with a tanh non-linearity. For `SELECTCOLUMN`, we compute

$$\tilde{\lambda}_i = \frac{\mathbf{h}_t W_Q^{\text{sc}} (y_i W_K^{\text{sc}})^T}{\sqrt{d_x}}; \lambda_i = \frac{\exp(\tilde{\lambda}_i)}{\sum_{j=1}^{|y|} \exp(\tilde{\lambda}_j)}; \Pr(a_t = \text{SELECTCOLUMN}[i] | a_{<t}, y) = \sum_{j=1}^{|y|} \lambda_j L_{j,i}^{\text{col}}$$

and similarly for `SELECTTABLE`.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

We implemented our model using PyTorch (Paszke et al., 2017). During preprocessing, the input of questions, column names and table names are tokenized and lemmatized with the StanfordNLP toolkit (Manning et al., 2014). Within the encoder, we use GloVe (Pennington et al., 2014) word embeddings, held fixed in training except for the 50 most common words in the training set. All word embeddings have dimension 300. The bidirectional LSTMs have hidden size 128 per direction, and use the recurrent dropout method of Gal and Ghahramani (2016) with rate 0.2. We stack 8

Table 2: Our main results (all numbers are exact match %).

(a) Accuracy on the Spider development and test sets, compared to the other approaches at the top of the dataset leaderboard as of Sept 24, 2019. The test set results were scored using the Spider evaluation server.

Model	Dev	Test
IRNet (Guo et al. (2019))	53.2	46.7
Global-GNN (Bogin et al. (2019a))	52.7	47.4
TPNet (anonymous)	55.4	48.5
RAT-SQL (ours)	60.6	53.7
<i>BERT</i>		
EditSQL + BERT (Zhang et al. (2019))	57.6	53.4
IRNet + BERT (Guo et al. (2019))	61.9	54.7
GIRN + BERT (anonymous)	60.2	54.8
TPNet + BERT (anonymous)	63.9	55.0

(b) Accuracy on the Spider development and test sets, by difficulty as defined by Yu et al. (2018c).

Split	Easy	Medium	Hard	Extra Hard	All
Dev	80.0	61.4	50.6	40.6	60.6
Test	73.1	60.1	45.3	24.8	53.7

(c) Accuracy (and $\pm 95\%$ confidence interval) of RAT-SQL ablations on the dev set. Schema linking makes a statistically significant difference ($p < 0.001$).

Model	Accuracy
RAT-SQL	58.52 \pm 0.84
RAT-SQL w/o alignment loss	58.61 \pm 0.59
RAT-SQL w/o schema linking relations	46.16 \pm 1.33

relation-aware self-attention layers on top of the bidirectional LSTMs. Within the relation-aware self-attention layers, we set $d_x = d_z = 256$, $H = 8$, and use dropout with rate 0.1. The position-wise feed-forward network has inner layer dimension 1024. Inside the decoder, we use rule embeddings of size 128, node type embeddings of size 64, and a hidden size of 512 inside the LSTM with dropout rate 0.21.

We used the Adam optimizer (Kingma and Ba, 2014) with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-9}$, which are defaults in PyTorch. During the first $warmup_steps = max_steps/20$ steps of training, we linearly increase the learning rate from 0 to 7.4×10^{-4} . Afterwards, the learning rate is annealed to 0, with formula $10^{-3} \left(1 - \frac{step - warmup_steps}{max_steps - warmup_steps}\right)^{-0.5}$. For all parameters, we used the default initialization method in PyTorch. We use a batch size of 20 and train for up to 40,000 steps.

4.2 DATASET AND METRICS

We use the Spider dataset (Yu et al., 2018b) for all our experiments. As described by Yu et al. (2018b), the training data contains 8,659 examples, including 1,659 examples (questions and queries, with the accompanying schemas) from the Restaurants (Popescu et al., 2003; Tang and Mooney, 2000), GeoQuery (Zelle and Mooney, 1996), Scholar (Iyer et al., 2017), Academic (Li and Jagadish, 2014), Yelp and IMDB (Yahmazadeh et al., 2017) datasets.

As Yu et al. (2018b) make the test set accessible only through an evaluation server, we perform most evaluations (other than the final accuracy measurement) using the development set. It contains 1,034 examples, with databases and schemas distinct from those in the training set. We report results using the same metrics as Yu et al. (2018a): exact match accuracy on all examples, as well as divided by difficulty levels specified in the dataset. As in previous work, these metrics do not measure the model’s performance on generating values within the queries.

4.3 RESULTS

In Table 2a we show accuracy on the (hidden) test set for RAT-SQL and compare to all other approaches that are at or near state-of-the-art (according to the official dataset leaderboard). RAT-SQL outperforms all other methods that, like RAT-SQL, are not augmented with BERT embeddings. It even comes within 1.3% of beating the best BERT-augmented model. Since the typical improvement achieved by BERT augmentation is about 7% for all models, we are hopeful that adding such augmentation to RAT-SQL will also lead to state-of-the-art performance among BERT models.

We also provide a breakdown of the accuracy by difficulty in Table 2b. As expected, performance drops with increasing difficulty. The overall generalization gap between development and test was strongly affected by the significant drop in accuracy (15%) on the extra hard questions.

Schema Linking Table 2c shows an ablation study without RAT-based schema linking relations. Schema linking makes a statistically significant improvement to accuracy ($p < 0.001$). The full




Figure 4: Alignment between the question “For the cars with 4 cylinders, which model has the largest horsepower?” and the database `car_1` schema (columns and tables).

model accuracy here differs from Table 2a because the latter shows the best single model from a hyper-parameter sweep (submitted for test evaluation) and the former gives the mean over ten runs.

Alignment Recall from Section 3 that we explicitly represent the alignment between question words and table columns which is used during decoding for column selection. The existence of the alignment matrix provides a mechanism for the model to align words to columns, but the additional terms in the loss encourage it to actually act like an alignment.

In our final model, the alignment loss terms do not make a difference in overall accuracy. This is surprising to us because in earlier development, the alignment loss did improve the model (statistically significantly, from 53.0% to 55.4%). We hypothesize that hyper-parameter tuning that caused us to increase encoding depth also eliminated the need for explicit supervision of alignment.

An accurate alignment representation has other benefits as well, such as identifying question words to copy when a constant is needed (not part of the Spider dataset evaluation). In Figure 4 we show the alignment generated by our model on an example from the development set.³ For the three key words that reference columns (“cylinders”, “model”, “horsepower”), the alignment matrix correctly identifies their corresponding column (`cylinders`, `model`, `horsepower`) and the table (`cars_data`) except it mistakenly aligns “model” to `cars_data` also instead of to `car_names`. The word “cars” aligns to the primary key of the `cars_data` table.

5 CONCLUSION

Despite the abundance of research in semantic parsing of text to SQL, many contemporary models struggle to learn good representations for a given database schema as well as to properly link column/table references in the question. These problems are related: to encode & use columns/tables from the schema, the model must reason about their role in the context of a given question. In this work, we present a unified framework for addressing the schema encoding and linking challenges. Thanks to relation-aware self-attention, it jointly learns schema and question word representations based on their alignment with each other and predefined schema relations.

Empirically, the RAT framework allows us to gain significant state of the art improvement on text-to-SQL parsing. Qualitatively, it provides a way to combine predefined *hard* schema relations and inferred *soft* self-attended relations in the same encoder architecture. We foresee this joint representation learning being beneficial in many learning tasks beyond text-to-SQL, as long as the input has predefined structure.

³ The full alignment also maps from column and table names, but those end up simply aligning to themselves or the table they belong to, so we omit them for clarity of presentation.

REFERENCES

- Ben Bogin, Matt Gardner, and Jonathan Berant. Global reasoning over database structures for text-to-sql parsing. *arXiv preprint arXiv:1908.11214*, 2019a.
- Ben Bogin, Matt Gardner, and Jonathan Berant. Representing schema structure with graph neural networks for text-to-sql parsing. *arXiv preprint arXiv:1905.06241*, 2019b.
- Catherine Finegan-Dollak, Jonathan K. Kummerfeld, Li Zhang, Karthik Ramanathan, Sesh Sadasivam, Rui Zhang, and Dragomir Radev. 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. Association for Computational Linguistics, 2018. URL <http://aclweb.org/anthology/P18-1033>.
- Yarin Gal and Zoubin Ghahramani. A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 29, pages 1019–1027. Curran Associates, Inc., 2016. URL <http://papers.nips.cc/paper/6241-a-theoretically-grounded-application-of-dropout-in-recurrent-neural-networks.pdf>.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. Towards complex text-to-sql in cross-domain database with intermediate representation. *arXiv preprint arXiv:1905.08205*, 2019.
- Pengcheng He, Yi Mao, Kaushik Chakrabarti, and Weizhu Chen. X-sql: reinforce schema representation with context. *arXiv preprint arXiv:1908.08113*, 2019.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 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, 2017. URL <http://www.aclweb.org/anthology/P17-1089>.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*, December 2014. URL <http://arxiv.org/abs/1412.6980>.
- Fei Li and H. V. Jagadish. Constructing an interactive natural language interface for relational databases. *Proceedings of the VLDB Endowment*, 8(1):73–84, September 2014. URL <http://dx.doi.org/10.14778/2735461.2735468>.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. The Stanford CoreNLP natural language processing toolkit. In *Association for Computational Linguistics (ACL) System Demonstrations*, pages 55–60, 2014. URL <http://www.aclweb.org/anthology/P/P14/P14-5010>.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. October 2017. URL <https://openreview.net/forum?id=BJJsrmfCZ>.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October 2014. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/D14-1162>.
- Ana-Maria Popescu, Oren Etzioni, , and Henry Kautz. Towards a theory of natural language interfaces to databases. In *Proceedings of the 8th International Conference on Intelligent User Interfaces*, pages 149–157, 2003. URL <http://doi.acm.org/10.1145/604045.604070>.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-Attention with Relative Position Representations. 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 464–468. Association for Computational Linguistics, 2018. doi: 10.18653/v1/N18-2074. URL <http://aclweb.org/anthology/N18-2074>.

- Lappoon R. Tang and Raymond J. Mooney. Automated construction of database interfaces: Integrating 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, 2000. URL <http://www.aclweb.org/anthology/W00-1317>.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Navid Yaghmazadeh, Yuepeng Wang, Isil Dillig, and Thomas Dillig. Sqlizer: Query synthesis from natural language. In *International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, ACM, pages 63:1–63:26, October 2017. URL <http://doi.org/10.1145/3133887>.
- Pengcheng Yin and Graham Neubig. A Syntactic Neural Model for General-Purpose Code Generation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 440–450. Association for Computational Linguistics, 2017. doi: 10.18653/v1/P17-1041. URL <http://aclweb.org/anthology/P17-1041>.
- Tao Yu, Michihiro Yasunaga, Kai Yang, Rui Zhang, Dongxu Wang, Zifan Li, and Dragomir Radev. SyntaxSQLNet: Syntax Tree Networks for Complex and Cross-Domain Text-to-SQL Task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1653–1663. Association for Computational Linguistics, 2018a. URL <http://aclweb.org/anthology/D18-1193>.
- 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. 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, 2018b. URL <http://aclweb.org/anthology/D18-1425>.
- 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. 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. Association for Computational Linguistics, 2018c. URL <http://aclweb.org/anthology/D18-1425>.
- John M. Zelle and Raymond J. Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2*, pages 1050–1055, 1996. URL <http://dl.acm.org/citation.cfm?id=1864519.1864543>.
- Rui Zhang, Tao Yu, He Yang Er, Sungrok Shim, Eric Xue, Xi Victoria Lin, Tianze Shi, Caiming Xiong, Richard Socher, and Dragomir Radev. Editing-based sql query generation for cross-domain context-dependent questions. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2019.
- Victor Zhong, Caiming Xiong, and Richard Socher. Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning. *arXiv:1709.00103 [cs]*, August 2017. URL <http://arxiv.org/abs/1709.00103>.

Table 3: Accuracy (exact match %) on development set with an oracle providing correct columns and tables (Oracle cols) and/or the AST sketch structure (Oracle sketch).

Model	Accuracy
RAT-SQL	60.6
RAT-SQL + Oracle cols	67.6
RAT-SQL + Oracle sketch	70.9
RAT-SQL + Oracle sketch + Oracle cols	99.4

A THE NEED FOR SCHEMA LINKING

One natural question is how often does the decoder fail to select the correct column, even with the schema encoding and linking improvements we have made. To answer this, we conducted an oracle experiment (see Table 3).

For "oracle sketch", at every grammar nonterminal the decoder is forced to make the correct choice so the final SQL sketch exactly matches that of the correct answer. The rest of the decoding proceeds as if the decoder had made the choice on its own. Similarly, "oracle cols" forces the decoder to output the correct column or table at terminal productions.

With both oracles, we see an accuracy of 99.4% which just verifies that our grammar is sufficient to answer nearly every question in the data set. With just "oracle sketch", the accuracy is only 70.9%, which means 73.5% of the questions that RAT-SQL gets wrong and could get right have incorrect column or table selection. Similarly, with just "oracle cols", the accuracy is 67.6%, which means that 82.0% of the questions that RAT-SQL gets wrong have incorrect structure. In other words, most questions have both column and structure wrong, so both problems will continue to be important to work on for the future.