

IGSQL: Database Schema Interaction Graph Based Neural Model for Context-Dependent Text-to-SQL Generation

Yitao Cai and Xiaojun Wan

Wangxuan Institute of Computer Technology, Peking University

Center for Data Science, Peking University

The MOE Key Laboratory of Computational Linguistics, Peking University

{caiyitao, wanxiaojun}@pku.edu.cn

Abstract

Context-dependent text-to-SQL task has drawn much attention in recent years. Previous models on context-dependent text-to-SQL task only concentrate on utilizing historical user inputs. In this work, in addition to using encoders to capture historical information of user inputs, we propose a database schema interaction graph encoder to utilize historical information of database schema items. In decoding phase, we introduce a gate mechanism to weigh the importance of different vocabularies and then make the prediction of SQL tokens. We evaluate our model on the benchmark SParC and CoSQL datasets, which are two large complex context-dependent cross-domain text-to-SQL datasets. Our model outperforms previous state-of-the-art model by a large margin and achieves new state-of-the-art results on the two datasets. The comparison and ablation results demonstrate the efficacy of our model and the usefulness of the database schema interaction graph encoder.

1 Introduction

The Text-to-SQL task aims to translate natural language texts into SQL queries. Users who do not understand SQL grammars can benefit from this task and acquire information from databases by just inputting natural language texts. Previous works (Li and Jagadish, 2014; Xu et al., 2017; Yu et al., 2018a; Bogin et al., 2019b; Huo et al., 2019) focus on context-independent text-to-SQL generation. However, in practice, users usually interact with systems for several turns to acquire information, which extends the text-to-SQL task to the context-dependent text-to-SQL task in a conversational scenario. Throughout the interaction, user inputs may omit some information that appeared before. This phenomenon brings difficulty for context-dependent text-to-SQL task.

Recently, context-dependent text-to-SQL task has attracted more attention. Suhr et al. (2018) conduct experiments on ATIS dataset (Dahl et al., 1994). Besides, two cross-domain context-dependent datasets SParC (Yu et al., 2019b) and CoSQL (Yu et al., 2019a) are released. Cross-domain means databases in test set differ from that in training set, which is more challenging.

EditSQL (Zhang et al., 2019) is the previous state-of-the-art model on SParC and CoSQL datasets and it focuses on taking advantages of previous utterance texts and previously predicted query to predict the query for current turn. Table 1 shows the user inputs, ground truth queries and predicted queries of EditSQL for an interaction. In the second turn, EditSQL views “Kacey” as the name of a dog owner. However, since the context of the interaction is about dogs, “Kacey” should be the name of a dog. This example shows that a model using only historical information of user inputs may fail to keep context consistency and maintain thematic relations.

According to (Yu et al., 2019b) and (Yu et al., 2019a), to maintain thematic relations, users may change constraints, ask for different attributes for the same topic when they ask the next questions. Thus, database schema items (i.e., *table.column*) in current turn should have relation with items in previous turn. For example, in Table 1, the second question x^2 adds a constraint of the name and asks for the age of a dog instead of the numbers of all dogs. The corresponding database schema items *Dogs.age* and *Dogs.name* in y^2 belong to the same table as *Dogs.** in previous query y^1 . Therefore, we propose to take historical information about database schema items into consideration.

In particular, we first construct a graph based on corresponding database, where graph nodes are database schema items and graph edges are primary-foreign keys and column affiliation. Short

x^1	how many dogs on the table
\tilde{y}^1	SELECT count (*) FROM Dogs
y^1	SELECT count (*) FROM Dogs
x^2	what is the age of Kacey
\tilde{y}^2	SELECT T2.age FROM owners as T1 JOIN Dogs AS T2 ON T1.owner_id = T2.owner_id WHERE T1.first_name = 1
y^2	SELECT age FROM dogs WHERE name = "Kacey"
x^3	which dog is highest weight on table – Do you want the name of the dog with the highest weight? – exactly
\tilde{y}^3	SELECT name FROM dogs ORDER BY weight DESC limit 1
y^3	SELECT name FROM dogs ORDER BY weight DESC limit 1
x^4	What is the size code of BUL – Did you mean the size code of dogs with a breed code BUL? – exactly
\tilde{y}^4	SELECT size_code FROM dogs WHERE breed_code = 1
y^4	SELECT size_code FROM dogs WHERE breed_code = "BUL"

Table 1: An example interaction. x^i is the input sequence in i -th turn and y^i is the corresponding ground truth query. \tilde{y}^i means that query is predicted by a model, which is EditSQL here.

distance between graph nodes appearing in previous query and current query can reveal the context consistency since there is usually an edge between the different attributes of the same topic. We then propose a database schema interaction graph encoder to model database schema items together with historical items. Empirical results on two large cross-domain context-dependent text-to-SQL datasets - SPaC and CoSQL show that our schema interaction graph encoder contributes to modeling context consistency and our proposed model with database schema interaction graph encoder substantially outperforms the state-of-the-art model.

Our main contributions are summarized as follows:

- Previous models failed to keep context consistency and predict queries in a conversation scenario. To remedy this, we propose a database schema interaction graph encoder for database schema encoding and it can keep context consistency for the context-dependent text-to-SQL task. Our implementations are

public available ¹.

- Our model with the database schema interaction graph encoder achieves new state-of-the-art performances on development and test sets of two cross-domain context-dependent text-to-SQL datasets, SPaC and CoSQL.

2 Related Work

Many studies have focused on context-independent text-to-SQL task. [Zhong et al. \(2017\)](#) split the vocabulary and use reinforcement learning. [Xu et al. \(2017\)](#) propose a sketched-based model, which decomposes the token prediction process into SELECT-clause prediction and WHERE-clause prediction, aiming at taking previous predictions into consideration. [Yu et al. \(2018a\)](#) further employ a tree-based SQL decoder so as to decode SQL queries with the help of SQL grammar. In order to encode database schemas, schemas are regarded as graphs and graph neural networks have been applied ([Bogin et al., 2019a,b](#)). [Guo et al. \(2019\)](#) design an intermediate representation to bridge the gap between natural language texts and SQL queries. [Choi et al. \(2020\)](#) utilize a sketch-based slot filling approach to synthesize SQL queries. [Wang et al. \(2019\)](#) attempt to align the database columns and their mentions in user inputs by using a relation-aware self attention.

Recently, context-dependent text-to-SQL task has drawn people’s attention. In-domain context-dependent benchmarks ATIS ([Suhr et al., 2018](#)) have been proposed. For ATIS, [Suhr et al. \(2018\)](#) utilize a sequence to sequence framework. Besides, they introduce an interaction-level encoder for incorporating historical user inputs and a segment copy mechanism to reduce the length of generation. Later, two large and complex cross-domain context-dependent dataset SPaC ([Yu et al., 2019b](#)) and CoSQL ([Yu et al., 2019a](#)) are proposed. In order to tackle cross-domain context-dependent text-to-SQL task, [Zhang et al. \(2019\)](#) propose the EditSQL model in order to capture features from historical user inputs, variant database schemas and previously predicted SQL query. EditSQL achieves the state-of-the-art performance on the two cross-domain datasets. Compared to EditSQL, our work further explore a new way to employ historical information of database schemas.

¹<https://github.com/headacheboy/IGSQL>

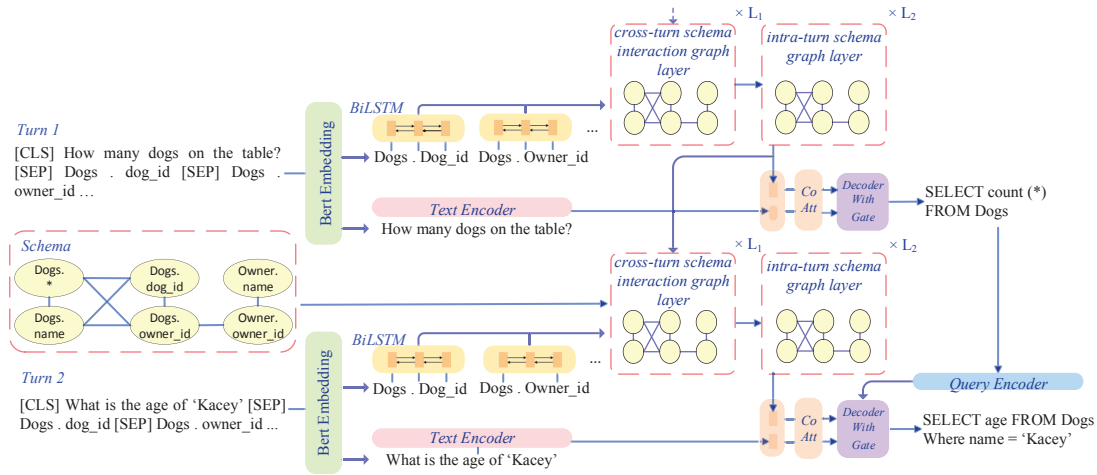


Figure 1: Overview of our IGSQl model. Modules with the same color share the same parameters.

3 Problem Setup

We define X as a series of natural language utterances of an interaction (i.e., user inputs), Y as corresponding ground-truth SQL queries, S as the set of database schema items (*table.column*) and R as the set of relations between schema items (primary-foreign keys and column affiliation). Let $X = \{x^1, x^2, \dots, x^{|X|}\}$, where $|X|$ is the number of utterances. x^i is the i -th utterance and x_j^i is the j -th token of it. y^i is the i -th SQL query corresponding to x^i and y_j^i is the j -th token of y^i . S consists of schema items $\{S^1, \dots, S^{|S|}\}$, where $|S|$ is the number of database schema items. At turn i , the model should make use of current and previous utterances $\{x^1, x^2, \dots, x^i\}$, database schema items S and their relations R to predict a SQL query \tilde{y}^i . The objective of the model is to maximize the probability of $\prod_{i=1}^{|I|} P(y^i | x^1, x^2, \dots, x^i)$.

4 IGSQl Model

Our model adopts an encoder-decoder framework with attention mechanism. Figure 1 shows the architecture of our model. The model has four main components: (1) a database schema interaction graph encoder, which consists of cross-turn schema interaction graph layers and intra-turn schema graph layers, (2) a text encoder that captures historical information of user inputs, (3) a co-attention module that updates outputs of text encoder and database schema interaction graph encoder, and (4) a decoder with a gated mechanism to weight the importance of different vocabularies. In addition, the model also uses BERT embedding.

We will first introduce the BERT embedding

in Section 4.1, and then introduce our database schema interaction graph encoder in Section 4.2, text encoder and co-attention module in Section 4.3 and decoder in Section 4.4.

4.1 BERT Embedding

BERT (Devlin et al., 2019) is a pre-trained language model. Employing BERT output as embeddings of user inputs and database schema items has proved effective in context-dependent text-to-SQL task (Hwang et al., 2019; Guo et al., 2019; Wang et al., 2019; Choi et al., 2020). Therefore, we leverage BERT to get the embeddings of user inputs and database schema items as other context-dependent text-to-SQL models do. We concatenate user inputs and database schema items by separating with a “[SEP]” token following (Hwang et al., 2019). The output of BERT model is used as the embeddings of user inputs and schema items.

4.2 Database Schema Interaction Graph Encoder

As shown in Table 1, previous model mistakes “Kacey” as the name of a dog owner. However, the interaction is all about dogs and “Kacey” should be the name of a dog. It shows that previous model does not perform well in modeling context consistency of an interaction.

For two database schema items appearing in two adjacent turns, short distance of items in the graph can reveal the context consistency. For example, the distance between *Dogs.**² and correct item *Dogs.name* is 1. Distance between *Dogs.** and wrong item *owners.name* is 3.

²*table.** is considered a special column in table.

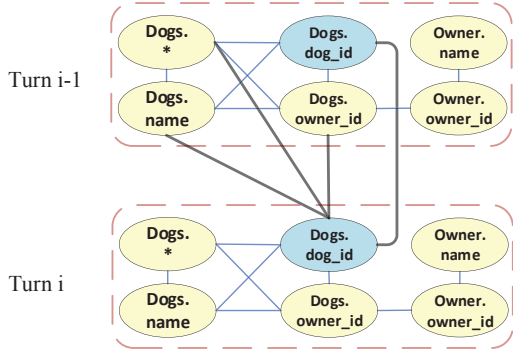


Figure 2: Database Schema Interaction Graph. We add black edges into the graph when we want to update the representation of the blue node (*Dogs.dog_id*) at turn i .

Therefore, we propose a database schema interaction graph encoder based on the database schema graph, attempting to model context consistency by using historical schema representations. The database schema interaction graph encoder consists of L_1 cross-turn schema interaction graph layers and L_2 intra-turn schema graph layers (L_1 and L_2 are hyper-parameters). Cross-turn schema interaction graph layers update schema item representations by using that in previous turn. Intra-turn schema graph layers further aggregate adjacent item representations in the same turn.

4.2.1 Graph Construction and Schema Items Encoding

We first introduce how we construct a graph based on database schema. We use database schema items as nodes. Each node has an edge linking to itself. There is an undirected edge between node t and node j according to relation set R if one of the following condition is satisfied: 1) node t and node j are the foreign-primary key pair; 2) node t and node j belong to the same table. We define the edge set as E .

A schema item *table.column* is divided into “*table*”, “*.*” and “*column*”. We use a BiLSTM with BERT embedding to encode tokens and average hidden state vectors of BiLSTM as the embedding of the schema item. The embedding of the j -th schema item at i -th turn is noted as r_j^i .

4.2.2 Cross-turn Schema Interaction Graph Layer

Figure 2 shows an example of the database schema interaction graph. The graph only allows node t in previous turn to update node j in current turn, when the distance between node t and node j in the orig-

inal graph constructed in Section 4.2.1 is less than or equal to 1. For example, if we want to update the representation of *Dogs.dog_id* at turn i , we add edges linking *Dogs.**, *Dogs.name*, *Dogs.owner_id* and *Dogs.dog_id* at turn $i - 1$ to *Dogs.dog_id* at turn i .

Note that we have L_1 cross-turn schema interaction graph layers for turn i . At the l -th layer, we obtain updated representation $z_t^{i,l}$ of the t -th schema item by using attention on outputs of the L_2 intra-turn schema graph layers at previous turn $\{g_t^{i-1,L_2}\}_{t=1}^{|S|}$ (which will be introduced in next subsection) and representations of previous layer $\{z_t^{i,l-1}\}_{t=1}^{|S|}$. We use item embedding r_t^i as the initial representation $z_t^{i,0}$. For simplicity, we omit turn index i and layer index l in the formulas of attention mechanism except the input $z_t^{i,l-1}$, g_t^{i-1,L_2} and output $z_t^{i,l}$.

At the l -th layer, we first use a feed-forward neural network with leakyReLU activation function for non-linear transformation. We use FFN to denote the feed-forward neural network with leakyReLU activation function.

$$\begin{aligned} u_t &= \text{FFN}(z_t^{i,l}) \\ \hat{u}_t &= \text{FFN}(g_t^{i-1,L_2}) \end{aligned} \quad (1)$$

We then apply attention mechanism as follows.

$$\begin{aligned} \xi_{t,j} &= \begin{cases} (u_t)^T W_1 u_j / \sqrt{d_1}, & [t, j] \in E \\ -\infty, & [t, j] \notin E \end{cases} \\ \hat{\xi}_{t,j} &= \begin{cases} (u_t)^T W_2 \hat{u}_j / \sqrt{d_1}, & [t, j] \in E \\ -\infty, & [t, j] \notin E \end{cases} \\ \alpha_{t,j} &= \frac{\exp(\xi_{t,j})}{\sum_v \exp(\xi_{t,v}) + \sum_k \exp(\hat{\xi}_{t,k})} \\ \hat{\alpha}_{t,j} &= \frac{\exp(\hat{\xi}_{t,j})}{\sum_v \exp(\xi_{t,v}) + \sum_k \exp(\hat{\xi}_{t,k})} \\ \tilde{u}_t &= \sum_j \alpha_{t,j} u_j + \sum_j \hat{\alpha}_{t,j} \hat{u}_j \end{aligned} \quad (2)$$

where d_1 is the dimension of u_t . W_1 and W_2 are weight matrices. $\alpha_{t,j}$ and $\hat{\alpha}_{t,j}$ are the attention scores. \tilde{u}_t is the t -th output vector of attention.

Following (Vaswani et al., 2017; Veličković et al., 2017), we extend attention mechanism to multi-head attention. We also add a sub-layer of feed-forward neural network with residual connec-

tion as in Transformer.

$$z_t^{i,l} = z_t^{i,l-1} + \text{FFN}(z_t^{i,l-1} + \tilde{u}_t) \quad (3)$$

where $z_t^{i,l}$ is the final output of layer l . There are L_1 cross-turn schema interaction graph layers and thus z_t^{i,L_1} is the final output of cross-turn schema interaction graph layers for the t -th schema item.

4.2.3 Intra-turn Schema Graph Layer

There are L_2 intra-turn schema graph layers following cross-turn schema interaction graph layers. In each intra-turn schema graph layer, we use almost the same attention mechanism as in the cross-turn schema interaction graph layer, except that we use the original graph constructed in Section 4.2.1. Since the original graph does not contain nodes in previous turn, the intra-turn schema graph layer can only update node representation by aggregating adjacent node representations in the same turn.

At each intra-turn schema graph layer layer l of turn i , it takes output vectors in previous layer $g_t^{i,l-1}$ as inputs and its output is $g_t^{i,l}$. $g_t^{i,0}$ is z_t^{i,L_1} . We then use attention mechanism to aggregate information. We also add a sub-layer of FFN and residual connection. For simplicity, we omit the turn index i and layer index l in attention except input $g_t^{i,l-1}$ and output $g_t^{i,l}$.

$$\begin{aligned} \mu_t &= \text{FFN}(g_t^{i,l-1}) \\ \tau_{t,j} &= \begin{cases} (\mu_t)^T W_3 \mu_j / \sqrt{d_2}, & [t, j] \in E \\ -\infty, & [t, j] \notin E \end{cases} \\ \beta_{t,j} &= \frac{\exp(\tau_{t,j})}{\sum_k \exp(\tau_{t,k})} \\ \tilde{\mu}_t &= \sum_j \beta_{t,j} \mu_j \\ g_t^{i,l} &= g_t^{i,l-1} + \text{FFN}(g_t^{i,l-1} + \tilde{\mu}_t) \end{aligned} \quad (4)$$

where W_3 is a weight matrix and d_2 is the dimension of μ_t . $\beta_{t,j}$ is the attention score of the j -th node to the t -th node. $\tilde{\mu}_t$ is the attention output. $g_t^{i,l}$ is the output of t -th schema item at layer l of turn i . Besides, We also extend attention to multi-head attention.

The final output of intra-turn schema graph layers for the t -th schema item is g_t^{i,L_2} .

4.3 Text Encoder and Co-Attention Module

We use a BiLSTM to encode tokens of an utterance text with BERT embedding. In order to capture

interaction history, we add an LSTM as interaction encoder and utilize turn-level attention, following (Zhang et al., 2019). The final representation of the t -th token in utterance i is denoted as h_t^i .

We also add a co-attention module between text tokens and schema items following (Zhang et al., 2019). The schema item vector \tilde{g}_t^i used in decoding phase is the concatenation of g_t^{i,L_2} and its corresponding attention vector over text. The representation of input text tokens \tilde{h}_t^i used in decoding phase is the concatenation of h_t^i and its corresponding attention vector over schema items. Due to page limit, we omit the details here, which can be found in (Zhang et al., 2019).

4.4 Decoder

In decoding phase, we first encode previously predicted query with a BiLSTM. We then exploit a LSTM decoder with attention (Bahdanau et al., 2015) to capture features from input text's token vectors, schema item vectors and previously predicted SQL query vectors. At j -th time step, We use attention on text token's vector \tilde{h}_t^i , database schema vector \tilde{g}_t^i and previously predicted SQL token's vector q_t . We thus get three context vectors. The final context vector c_j is the concatenation of these three context vectors.

We follow (Suhr et al., 2018) to make prediction of SQL tokens based on SQL reserved words, database schema items and previous predicted SQL tokens. We also add a gate mechanism to introduce the importance of these three vocabularies. For simplicity, we omit turn index i in decoder step except \tilde{y}_j^i .

The gate mechanism is introduced to measure the importance of three vocabularies.

$$\begin{aligned} \tilde{o}_j &= \tanh(W_o([o_j; c_j] + b_o)) \\ \zeta_m &= \sigma(W_m \tilde{o}_j + b_m) \\ m &\in \{res, sch, que\} \end{aligned} \quad (5)$$

where o_j is the j -th hidden vector of the LSTM decoder. c_j is the context vector. $[\cdot]$ is the concatenation operator and \tilde{o}_j is the non-linear transformation of $[o_j; c_j]$. σ is the sigmoid function. res, sch, que represent SQL reserved words, database schema items and previously predicted SQL tokens respectively and $\zeta_{res}, \zeta_{sch}, \zeta_{que}$ represent the importance of these three kinds of tokens.

	SParC	CoSQL
cross-domain	✓	✓
Interaction	4298	3007
Train	3034	2164
Dev	422	292
Test	842	551
User Questions	12726	15598
Databases	200	200
Tables	1020	1020
Vocab	3794	9585
Avg Turn	3.0	5.2

Table 2: Statistics of SParC, CoSQL

We then predict SQL tokens as follows.

$$\begin{aligned}
 p_1(\tilde{y}_j^i = w) &= \frac{1}{Z} \exp(\zeta_{res} \cdot \mathbf{w}^T (W_{res} \tilde{o}_j + b_{res})) \\
 p_2(\tilde{y}_j^i = S_t) &= \frac{1}{Z} \exp(\zeta_{sch} \cdot (\tilde{g}_t^i W_{sch} \tilde{o}_j)) \\
 p_3(\tilde{y}_j^i = \tilde{y}_t^{i-1}) &= \frac{1}{Z} \exp(\zeta_{que} \cdot (q_t W_{que} \tilde{o}_j))
 \end{aligned} \tag{6}$$

where \mathbf{w} is the one-hot vector of word w . q_t and \tilde{g}_t^i are query vector and schema item vector that are mentioned before. The final generation probability $p(\tilde{y}_j^i)$ is $p_1(\tilde{y}_j^i) + p_2(\tilde{y}_j^i) + p_3(\tilde{y}_j^i)$. Z is the normalization factor that ensures $\sum_{v \in V} p(v)$ is 1, where V is the whole vocabulary. The loss function is $\sum_i \sum_j -\log(p(\tilde{y}_j^i))$

5 Implementation Details

We use Adam optimizer (Kingma and Ba, 2015) to optimize the loss function. The initial learning rate except BERT model is 1e-3, while the initial learning rate of BERT model is 1e-5. We use learning rate warmup over the first 1000 steps. The learning rate will be multiplied by 0.8 if the loss on development set increases and the token accuracy on development set decreases. The number of cross-turn schema interaction graph layer L_1 is 2, while the number of intra-turn schema graph layer L_2 is 1. The dimensions d_1 and d_2 are both 300. For encoder and decoder, the hidden size of the one layer LSTM and BiLSTM are 300. Besides, we use batch re-weighting to reweigh the loss function following (Suhr et al., 2018). For BERT embedding, following EditSQL, we use the pre-trained BERT base model in order to make fair comparison.

6 Experiments

6.1 Experiment Setup

Datasets. We conduct experiments on two large-scale cross-domain context-dependent SQL generation datasets, SParC (Yu et al., 2019b) and CoSQL (Yu et al., 2019a). In comparison with previous context-dependent dataset ATIS (Dahl et al., 1994), SParC and CoSQL are more complex since they contain more databases and adopt a cross-domain task setting, where the databases of training set differ from that of development set and test set. Statistics of SParC and CoSQL are shown in Table 2.

Evaluation Metrics. Yu et al. (2018b) introduce exact set match accuracy to replace string match accuracy by taking queries with same constraints but different orders as the same query. In SParC and CoSQL, we use question match accuracy and interaction match accuracy as evaluation metrics. Question match accuracy is the average exact set match accuracy over all questions, while interaction match accuracy is the average exact set match accuracy over all interactions.

Baseline Models. We compare our model with following baseline models.

- **Context dependent Seq2Seq (CD S2S).** This model is originated in (Suhr et al., 2018) for ATIS dataset. Yu et al. (2019b) adapt this model to cross-domain setting by adding a BiLSTM to encode schema items and modifying the decoder to generate different schema items according to databases.
- **SyntaxSQL-con.** This model is originated in (Yu et al., 2018a), which utilizes SQL grammars for decoder. Yu et al. (2019b) adapt this model to context-dependent setting by adding LSTM encoders to encode historical user inputs and historical SQL queries.
- **EditSQL.** The model is proposed by (Zhang et al., 2019). In addition to modules for encoding historical user inputs and corresponding SQL queries, it also contains a copy mechanism to copy tokens from previous SQL queries.

6.2 Experiment Results

Results of these baseline models and our proposed IGSQ model are shown in Table 3. Our model surpasses the previous state-of-the-art model EditSQL.

Method	SPaRC Dev		SPaRC Test		CoSQL Dev		CoSQL Test	
	Ques	Int	Ques	Int	Ques	Int	Ques	Int
CD S2S	21.9	8.1	23.2	7.5	13.8	2.1	13.9	2.6
SyntaxSQL-con	18.5	4.3	20.2	5.2	15.1	2.7	14.1	2.2
EditSQL*	47.2	29.5	47.9	25.3	39.9	12.3	40.8	13.7
IGSQL*	50.7	32.5	51.2	29.5	44.1	15.8	42.5	15.0

Table 3: Results of models in SPaRC and CoSQL datasets. Ques means question match accuracy. Int means interaction match accuracy. * means that results are enhanced by BERT embedding.

Turns	SPaRC		CoSQL	
	EditSQL	IGSQL	EditSQL	IGSQL
1	62.2	63.2	50.0	53.1
2	45.1	50.8	36.7	42.6
3	36.1	39.0	34.8	39.3
4	19.3	26.1	43.0	43.0
>4	0	0	23.9	31.0

Table 4: Exact match accuracy w.r.t. turn number on development sets.

Hardness	SPaRC		CoSQL	
	EditSQL	IGSQL	EditSQL	IGSQL
Easy	68.8	70.9	62.7	66.3
Medium	40.6	45.4	29.4	35.6
Hard	26.9	29.0	22.8	26.4
Extra	12.8	18.8	9.3	10.3

Table 5: Exact match accuracy w.r.t. different hardness level on development sets.

IGSQL achieves substantial improvement on question match accuracy by 3.5, 3.3 points on SPaRC development and test sets and 4.2, 1.7 points on CoSQL development and test sets, respectively. As for interaction match accuracy, IGSQL improves by 3, 4.2 points on SPaRC development and test sets, and 3.5, 1.3 points on CoSQL development and test sets. Results demonstrate the effectiveness of our model.

Table 4 shows the exact match accuracy of interaction with respect to different turn number. In both datasets, performances on interactions with one turn improve less. In SPaRC, performances on interactions with two turns and four turns improve the most, while in CoSQL, performances on interaction with two turns and larger than four turns improve the most. These results demonstrate that our database schema interaction graph encoder contributes to modeling schema items in conversational scenarios.

Table 5 lists the exact match accuracy with respect to different hardness level. Results in the

table show that performance at each hardness level improves. The results indicate that capturing historical database schema information can not only improve the accuracy of easy questions, but also answer harder questions more accurately.

6.3 Ablation Study

In order to verify the usefulness of our database schema interaction graph encoder, we conduct several ablation experiments as follows.

w/o cross-turn schema interaction graph layer.

In this experiment, we discard cross-turn schema interaction graph layers. In this setting, our model cannot encode historical database schema information.

w/o intra-turn schema graph layer. In this experiment, we discard intra-turn schema graph layers to examine whether these layers are useful.

GRU interaction layer. One of the most common way to employ historical information of database schema items is to update node representation directly from historical vector of the same node. For example, in Figure 2, we can use a GRU by taking representation of *Dogs.dog_id* at turn $i - 1$ and its BERT embedding at turn i as input. The output of GRU is the vector of *Dogs.dog_id* at turn i . In this experiment, we use a GRU to replace cross-turn schema interaction graph layers.

Fully-connected interaction layer. To examine the effectiveness of our design of schema interaction graph, we make experiment that replaces the schema interaction graph with fully connected graph. Taking Figure 2 as an example, to update representation of blue node at turn i , there are edges connecting blue node at turn i to all nodes at turn $i - 1$.

Since the test sets of SPaRC and CoSQL are not public, we carry out the ablation experiments only on development sets of these two datasets. Table 6 shows the results of ablation experiments. Our full model achieves about 2 points improvement compared with the model without cross-turn schema in-

Method	SPaC		CoSQL	
	Ques Match	Int Match	Ques Match	Int Match
IGSQL	50.7	32.5	44.1	15.8
w/o cross-turn schema interaction graph layer	47.6(-3.1)	29.5(-3.0)	41.9(-2.2)	14.0(-1.8)
w/o intra-turn schema graph layer	50.2(-0.5)	31.1(-1.4)	42.9(-1.2)	14.0(-1.8)
GRU interaction layer	48.2(-2.5)	29.2(-3.3)	41.0(-3.1)	14.1(-1.7)
Fully-connected interaction layer	48.2(-2.5)	29.0(-3.5)	42.0(-2.1)	13.0(-2.8)

Table 6: Ablation study on development sets. Numbers in brackets are performance differences compared to IGSQL.

x^1	Which cartoon aired first?
EditSQL	SELECT title FROM cartoon ORDER BY original_air_date LIMIT 1
IGSQL	SELECT title FROM cartoon ORDER BY original_air_date LIMIT 1
y^1	SELECT title FROM cartoon ORDER BY original_air_date LIMIT 1
x^2	What was the last cartoon to air?
EditSQL	SELECT T1.title FROM cartoon AS T1 JOIN tv_channel AS T2 ON T1.channel = T2.id JOIN tv_series AS T3 ON T2.id = T3.channel ORDER BY T3.air_date LIMIT 1
IGSQL	SELECT title FROM cartoon ORDER BY original_air_date DESC LIMIT 1
y^2	SELECT title FROM cartoon ORDER BY original_air_date DESC LIMIT 1
x^3	What channel was it on?
EditSQL	SELECT channel FROM tv_series ORDER BY air_date LIMIT 1
IGSQL	SELECT channel FROM cartoon ORDER BY original_air_date DESC LIMIT 1
y^3	SELECT channel FROM cartoon ORDER BY original_air_date DESC LIMIT 1
x^4	What is the production code?
EditSQL	select T1.production_code FROM cartoon AS T1 JOIN tv_channel AS T2 ON T1.channel = T2.id JOIN tv_series AS T3 ON T2.id = T3.channel ORDER BY T3.air_date LIMIT 1
IGSQL	SELECT production_code FROM cartoon ORDER BY original_air_date DESC LIMIT 1
y^4	SELECT production_code FROM cartoon ORDER BY original_air_date DESC LIMIT 1

Table 7: An example of an interaction in CoSQL. x^i is the input sequence at i -th turn and y^i is the corresponding ground truth query. We show the predictions of EditSQL and IGSQL and mark the differences with red color.

teraction graph layers and the model with GRU interaction layer. Besides, our model achieves about 1 point improvement compared with the model without intra-turn schema graph layers. These results indicate that our cross-turn and intra-turn schema graph layers are very helpful.

The difference between cross-turn schema interaction graph layer and fully-connected interaction layer is how we add edges between nodes at turn $i - 1$ and turn i . Compared to fully-connected interaction layer, the schema interaction graph introduces a distance restriction when adding edges. Our model with schema interaction graph performs substantially better, which shows that our design of schema interaction graph can significantly help our model to keep context consistency.

6.4 Case Study

In Table 7, we show an interaction with four turns. We also provide the predictions of EditSQL and IGSQL and mark the differences with

red color. After the first turn, EditSQL confuses *cartoon.original_air_date* with *tv_series.air_date*. Our proposed IGSQL model successfully obtains answers in the correct order by taking historical information of database schema items into account.

7 Conclusion and Future work

In this paper, we focus on context-dependent cross-domain SQL generation task. We find that previous state-of-the-art model only takes historical user inputs and previously predicted query into consideration, but ignores the historical information of database schema items. Thus we propose a model named IGSQL to model database schema items in a conversational scenario. Empirical results demonstrate the efficacy of our model. We also conduct ablation experiments to reveal the significance of our database schema interaction graph encoder. For future work, we will explore methods attempting to solve hard and extra hard questions.

Acknowledgments

This work was supported by National Natural Science Foundation of China (61772036), Beijing Academy of Artificial Intelligence (BAAI) and Key Laboratory of Science, Technology and Standard in Press Industry (Key Laboratory of Intelligent Press Media Technology). We appreciate the anonymous reviewers for their helpful comments. Xiaojun Wan is the corresponding author.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*.
- Ben Bogin, Jonathan Berant, and Matt Gardner. 2019a. [Representing schema structure with graph neural networks for text-to-SQL parsing](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4560–4565, Florence, Italy. Association for Computational Linguistics.
- Ben Bogin, Matt Gardner, and Jonathan Berant. 2019b. 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 3650–3655.
- 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](#).
- 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*, pages 43–48. Association for Computational Linguistics.
- 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.
- 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.
- Siyu Huo, Tengfei Ma, Jie Chen, Maria Chang, Lingfei Wu, and Michael J Witbrock. 2019. Graph enhanced cross-domain text-to-sql generation. In *Proceedings of the Thirteenth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-13)*, pages 159–163.
- Wonseok Hwang, Jinyeung Yim, Seunghyun Park, and Minjoon Seo. 2019. [A comprehensive exploration on wikisql with table-aware word contextualization](#). *CoRR*, abs/1902.01069.
- 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*.
- Fei Li and HV Jagadish. 2014. Constructing an interactive natural language interface for relational databases. *Proceedings of the VLDB Endowment*, 8(1):73–84.
- Alane Suhr, Srinivasan Iyer, and Yoav Artzi. 2018. Learning to map context-dependent sentences to executable formal queries. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2238–2249.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2017. Graph attention networks. *arXiv preprint arXiv:1710.10903*.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2019. [Rat-sql: Relation-aware schema encoding and linking for text-to-sql parsers](#).
- Xiaojun Xu, Chang Liu, and Dawn Song. 2017. [Sql-net: Generating structured queries from natural language without reinforcement learning](#). *CoRR*, abs/1711.04436.
- Tao Yu, Michihiro Yasunaga, Kai Yang, Rui Zhang, Dongxu Wang, Zifan Li, and Dragomir Radev. 2018a. 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.
- Tao Yu, Rui Zhang, Heyang Er, Suyi Li, Eric Xue, Bo Pang, Xi Victoria Lin, Yi Chern Tan, Tianze Shi, Zihan Li, et al. 2019a. Cosql: A conversational text-to-sql challenge towards cross-domain natural language interfaces to databases. 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 1962–1979.

Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 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.

Tao Yu, Rui Zhang, Michihiro Yasunaga, Yi Chern Tan, Xi Victoria Lin, Suyi Li, Heyang Er, Irene Li, Bo Pang, Tao Chen, et al. 2019b. Sparc: Cross-domain semantic parsing in context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4511–4523.

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 5341–5352, 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.