

# Exploring Underexplored Limitations of Cross-Domain Text-to-SQL Generalization

Yujian Gan<sup>1</sup> Xinyun Chen<sup>2</sup> Matthew Purver<sup>1,3</sup>

<sup>1</sup>Queen Mary University of London    <sup>2</sup>UC Berkeley    <sup>3</sup>Jožef Stefan Institute  
{y.gan,m.purver}@qmul.ac.uk    xinyun.chen@berkeley.edu

## Abstract

Recently, there has been significant progress in studying neural networks for translating text descriptions into SQL queries under the zero-shot cross-domain setting. Despite achieving good performance on some public benchmarks, we observe that existing text-to-SQL models do not generalize when facing domain knowledge that does not frequently appear in the training data, which may render the worse prediction performance for unseen domains. In this work, we investigate the robustness of text-to-SQL models when the questions require rarely observed domain knowledge. In particular, we define five types of domain knowledge and introduce Spider-DK (DK is the abbreviation of domain knowledge), a human-curated dataset based on the Spider benchmark for text-to-SQL translation. NL questions in Spider-DK are selected from Spider, and we modify some samples by adding domain knowledge that reflects real-world question paraphrases. We demonstrate that the prediction accuracy dramatically drops on samples that require such domain knowledge, even if the domain knowledge appears in the training set, and the model provides the correct predictions for related training samples. <sup>1</sup>

## 1 Introduction

Research on cross-domain text-to-SQL benchmarks has led to numerous advances. Recent works (Zhao et al., 2021; Rubin and Berant, 2021; Hui et al., 2021) have achieved over 70% accuracy on Spider benchmark (Yu et al., 2018) and over 90% accuracy on WikiSQL benchmark (Zhong et al., 2017), which seems to suggest that existing models already solved most problems in this field. However, the follow-up studies from Deng et al. (2021); Gan et al. (2021); Suhr et al. (2020); Shaw et al. (2021); Oren et al. (2020); Keysers et al. (2020)

show that the generalization performance is much worse in more challenging scenarios. For example, Deng et al. (2021) investigate the cases when the explicit mentions of database columns are removed from the question. Similarly, (Gan et al., 2021) observe that the model accuracy dramatically drops by replacing schema-related words with some synonyms. On the other hand, Suhr et al. (2020) find that the generalization to other databases is much worse, due to the distribution shift of both questions and SQL queries. These papers introduce important challenges for improving the generalization performance, i.e., the model trained on a cross-domain text-to-SQL dataset (e.g., Spider (Yu et al., 2018)) does not generalize to a new external database. However, the performance degradation is somehow expected for the following reasons. First, removing the explicit mentions breaks the assumptions that make the schema linking effective. Second, SQL queries in other databases could come from a different distribution; e.g., according to the hardness criteria defined by Spider benchmark, over 40% Spider SQL queries are Medium hardness, but there are less than 10% Medium SQL queries in the GeoQuery dataset (Zelle and Mooney, 1996).

In this work, we demonstrate that the generalization performance could be poor even when both the NL questions and SQL queries follow the similar distribution to the training set. Specifically, we constructed Spider-DK, a challenging variant of the Spider development set, with the focus of evaluating the model understanding of domain knowledge. A domain means a certain type of application scenarios; for example, the Spider benchmark includes various distinct domains such as geography and university. Cross-domain text-to-SQL research aims to build a text-to-SQL model that can generate correct SQL queries and generalize to different domains. Therefore, one main challenge of cross-domain text-to-SQL generalization is to understand different knowledge required by different domains.

<sup>1</sup>Our dataset is available at <https://github.com/ygan/Spider-DK>.

<b>T1</b>	SELECT Columns Mentioned by Omission
<b>NL</b>	<i>Find the <u>name</u> of the teacher who ...</i>
<b>SQL</b>	<code>select <u>firstname</u> , <u>lastname</u> from ...</code>
<b>T2</b>	Simple Inference Required
<b>NL</b>	<i>... order of their date of birth <u>from old to young</u>.</i>
<b>SQL</b>	<code>... order by date_of_birth <u>asc</u></code>
<b>T3</b>	Synonyms Substitution in Cell Value Word
<b>NL</b>	<i>List the state in the <u>US</u> ...</i>
<b>SQL</b>	<code>... where billing_country = "<u>USA</u>" ...</code>
<b>T4</b>	One Non-Cell Value Word Generate a Condition
<b>NL</b>	<i>How many students got <u>accepted</u> after the tryout?</i>
<b>SQL</b>	<code>... from tryout where <u>decision</u>="yes"</code>
<b>T5</b>	Easy to Conflict with other Domains
<b>NL</b>	<i>... with <u>max speed</u> higher than 1000.</i>
<b>SQL</b>	<code>... where <u>max_speed</u> &gt; 1000</code>

Table 1: Five types of domain knowledge extracted from Spider training set. We name them as T1 to T5.

For example, the university domain usually needs the knowledge of different job titles and genders, while the geography domain emphasizes more on the knowledge of places instead of people. We show that the state-of-the-art models consistently fail in cases when specific domain knowledge is required for prediction, even if the domain knowledge is moderately mentioned in the training data, and the models accurately predict the corresponding training samples. Such discrepancy suggests that the models do not properly learn the domain knowledge in order to fit the training set, thus improving the model capability to capture the domain knowledge is an important direction towards achieving the cross-domain generalization for text-to-SQL applications. To our knowledge, we are the first work investigating the text-to-SQL model capability of understanding the domain knowledge provided in the training set, and generalizing the knowledge to new problems.

## 2 Spider-DK Dataset

### 2.1 Overview

We construct the Spider-DK benchmark by selecting samples from the Spider development set that require domain knowledge understanding, and we also manually modify some samples to incorporate domain knowledge. The purpose of building Spider-DK is to simulate the scenario where specific domain knowledge is involved in the users’ utterance query. Domain knowledge is often used unnoticedly, which makes some domain knowledge

unavoidable. For example, in the T5 of Table 1, the direct use of the `max_speed` column annotation raises a domain knowledge problem. We discuss the details of this problem later in Section 2.2.

Spider-DK contains 535 NL-SQL pairs drawn from the Spider development set, where 270 pairs are the same as the original Spider samples, while the rest 265 pairs are modified to incorporate the domain knowledge. We categorize the types of domain knowledge required in Spider-DK, which makes it easy for breakdown analysis. Spider-DK is smaller than the Spider development set, because not every domain or example can be easily modified to incorporate some domain knowledge. Besides, it is hard to evaluate the model generalization ability for domain knowledge if keeping too many original Spider examples that do not require domain knowledge.

In particular, the distribution of the SQL query hardness in Spider-DK is close to the original Spider, i.e., easy accounts for 20.6%, medium accounts for 41.8%, hard accounts for 14.8%, and extra hard accounts for 19.1%<sup>2</sup>. We define five types of domain knowledge in Table 1. In Spider-DK, T1 accounts for 28.7% of samples, T2 accounts for 24.5%, T3 accounts for 27.5%, T4 accounts for 8.3%, and T5 accounts for 12.5%.

We curate the Spider-DK by modifying only questions or both questions and SQL queries, as shown in Table 2. We carefully add the domain knowledge into the utterance to ensure that the new utterance follows the domain knowledge required by existing Spider samples and does not raise ambiguity. Most domain knowledge in Spider-DK is similar to that in the Spider training set. Compared to the evaluation sets in (Suhr et al., 2020), Spider-DK is easier and closer to the training data and focuses only on domain knowledge, and we provide more discussion below.

### 2.2 Domain Knowledge

Different SQL databases could require very different domain knowledge. As shown in (Suhr et al., 2020), the state-of-the-art models on Spider achieve much worse performance on earlier SQL benchmarks such as ATIS and GeoQuery (Iyer et al., 2017; Zelle and Mooney, 1996). However, we argue that the failure of generalization is expected to some extent, because without seeing in-domain examples, some domain knowl-

<sup>2</sup>The Spider benchmark defines four hardness levels.

Only Modify the NL	
<b>Spider</b>	<i>... in the order of birth date.</i>
<b>Spider-DK</b>	<i>... order of their birth date from old to young.</i>
Modify both NL and SQL	
<b>Spider</b>	<i>Compute the average age of dogs.</i> select avg(age) from dogs
<b>Spider-DK</b>	<i>Compute the average age of <u>abandoned</u> dogs.</i> select avg(age) from dogs where <u>abandoned_y = 1</u>

Table 2: Examples of Spider question and/or SQL modifications made in Spider-DK.

edge required by these datasets is even hard to infer for experienced programmers. For example, we asked five computer science graduate students to write the SQL query for the question ‘how many major cities are there?’ in GeoQuery, but none of them gave the correct answer. This question requires the domain knowledge that major means ‘population > 150000’, which is hard to infer without looking at GeoQuery training set. Therefore, while acquiring general-purpose domain knowledge is also important, we believe that the failure of generalization to questions requiring similar domain knowledge to the training set could be more problematic, which motivates our design of Spider-DK benchmark.

We study five types of domain knowledge (name them as T1 to T5) shown in Table 1. T1 requires the models to understand that the user queries two columns by an omitted expression.

T2 requires the models to infer the correct queries, e.g., if the T2 utterance in Table 1 modified from ‘date of birth’ to ‘age’, the model should output desc not asc. Note that the Spider training set contains both ‘date of birth’ and ‘age’ along with ‘old to young’.

T3 requires the models to recognize the cell value synonym substitution. Some synonym substitutions base on their adjective form, such as ‘singer whose country is France’ and ‘French singer’.

Although the number of T4 is the least in Spider-DK, it is not uncommon in the Spider training set. Unlike the GeoQuery major example mentioned above, T4 only includes the conditions whose column type is similar to boolean. For example, in Table 1 and 2, the column decision only contain yes and no, while abandoned\_y only contain 1 and 0. Therefore, the key to solving T4 is whether the model can distinguish whether the column is a

boolean-like type, but the difficulty is that the word varies in different domains.

Although T5 seems simple and does not seem to contain domain knowledge, the models that generate SQL structure and schema items separately are easy to mispredict in T5. A review (Gan et al., 2020) shows that most models follow the separate generation pattern, i.e., these models may use the same word twice in both generating schema items and SQL structure. Because, in other domain training data, the models learn to generate a max() function when the utterance contains a word max. Therefore, these models may use the word max twice to generate the max(max\_speed) for T5 utterance instead of a simple max\_speed.

## 3 Experiments

### 3.1 Experimental Setup

We evaluate the previous state-of-the-art models on the Spider-DK and Spider (Yu et al., 2018). As discussed in Section 2.1, the Spider test set is not publicly accessible, and thus Spider-DK does not contain a test set. We extracted 535 examples corresponding to Spider-DK from Spider for evaluation instead of using a whole Spider development set for better comparison. In addition, we select 125 examples with domain knowledge from the training set to evaluate the training effect. Therefore, there are three evaluation sets:

- **Spider<sub>T</sub>**: 125 examples drawn from the Spider training set.
- **Spider<sub>D</sub>**: 535 examples drawn from the Spider development set.
- **Spider-DK**: Spider-DK development set with 535 examples.

We evaluate open-source models that reach competitive performance on Spider: GNN (Bogin et al., 2019), IRNet (Guo et al., 2019), RAT-SQL (Wang et al., 2020) with and without BERT (Devlin et al., 2019), and RAT-SQL + GAP (Shi et al., 2020). We present their results of the 265 Spider-DK domain knowledge examples and analyze their performance in each knowledge type. Our evaluation is based on the exact match metric defined in the original Spider benchmark, which measures whether the predicted query without condition values as a whole is equivalent to the gold query.

### 3.2 Main Results

Table 3 presents the exact match accuracy of different models on Spider<sub>T</sub>, Spider<sub>D</sub>, and Spider-DK.

model	Spider <sub>T</sub>	Spider <sub>D</sub>	Spider-DK
GNN (Bogin et al., 2019)	61.6%	46.2%	26.0%
IRNet (Guo et al., 2019)	87.2%	53.8%	33.1%
RAT-SQL (Wang et al., 2020)	93.6%	61.1%	35.8%
RAT-SQL + BERT (Wang et al., 2020)	92.0%	<b>73.3%</b>	40.9%
RAT-SQL + GAP (Shi et al., 2020)	<b>98.4%</b>	67.8%	<b>44.1%</b>

Table 3: Exact match accuracy on the Spider<sub>T</sub>, Spider<sub>D</sub> and Spider-DK, where models are trained on the original Spider training set.

Approach	ALL	T1	T2	T3	T4	T5
GNN	6.8%	5.3%	7.6%	2.7%	8.3%	21.2%
IRNet	19.2%	<b>9.2%</b>	4.6%	42.4%	4.5%	27.2%
RAT-SQL	16.6%	2.6%	13.8%	26.0%	<b>9.1%</b>	36.4%
RAT-SQL + BERT	19.6%	3.9%	12.3%	41.1%	4.5%	30.3%
RAT-SQL + GAP	<b>27.1%</b>	7.9%	<b>20.0%</b>	<b>53.4%</b>	<b>9.1%</b>	<b>42.4%</b>

Table 4: Break down exact match accuracy in the Spider-DK examples containing domain knowledge.

All models are trained on the original Spider training set. Compared to Spider<sub>D</sub>, the performance of all models has significantly dropped by about 20% to 30% on Spider-DK. Although the Spider-DK is designed based on the Spider<sub>T</sub>, whose exact match evaluation is pretty high, these models can not generalize to the Spider-DK well.

In particular, although RAT-SQL + BERT achieves better performance on Spider<sub>D</sub> than RAT-SQL + GAP, RAT-SQL + GAP outperforms RAT-SQL + BERT on Spider-DK, indicating that GAP facilitates the model to grasp a better understanding of domain knowledge. Despite some improvement achieved by recent models, the results show that domain knowledge understanding is still a considerable gap toward the realization of cross-domain text-to-SQL generation.

### 3.3 Performance on Knowledge Type Splits

To better understand the performance facing the domain knowledge, we present the breakdown accuracies of different domain knowledge types in Table 4. RAT-SQL + GAP unsurprisingly achieves the best performance on all examples and outperforms other models from T2 to T5. However, IRNet surprisingly obtains an overall accuracy close to the RAT-SQL + BERT, because IRNet integrates a ConceptNet (Speer and Havasi, 2012) to recognize the country, state, and city synonyms, which can improve its accuracy in T3. The GNN and RAT-SQL perform relatively poorly on T3 because they do not have extra knowledge components such as ConceptNet. Besides, GNN trains its embeddings from scratch, and RAT-SQL uses GLOVE (Pennington et al., 2014) that has been shown worse

(T1)NL	Find the <u>name</u> of the professionals ...
Pred	select <u>first_name</u> from ...
Gold	select <u>first_name</u> , <u>last_name</u> from ...
(T2)NL	... sorted from <u>oldest</u> to <u>youngest</u> ?
Pred	... order by birth_date <u>desc</u>
Gold	... order by birth_date <u>asc</u>
(T3)NL	List all <u>American</u> airline names and their abbreviations.
Pred	select airline, abbreviation from airlines
Gold	select airline, abbreviation from airlines where country = 'USA'
(T4)NL	What is the average age of all the <u>abandoned</u> dogs?
Pred	select avg(age) from dogs
Gold	select avg(age) from dogs where abandoned_y = 1
(T5)NL	Show <u>average</u> attendance for all stadiums ...
Pred	select <u>avg(average)</u> from stadium ...
Gold	select <u>average</u> from stadium ...

Table 5: Sample wrong predictions of RAT-SQL + GAP in each type of domain knowledge.

than BERT in many scenarios. Although ConceptNet helps IRNet in T3, it is not a generalization method for solving other domain knowledge problems. However, even the best-performing T3 is still far from the accuracy in Spider<sub>D</sub>, which shows that there is still much room for improvement.

### 3.4 Error Analysis

Table 5 presents five error examples in each knowledge type drawn from the prediction of RAT-SQL + GAP. These error predictions are similar to the training examples shown in Table 1. There are three reasons why existing models can not perform well in the Spider-DK. The first reason is that some domain knowledge is not common enough in the training set. For example, in T2, the phrase 'from old to young' appears more often with age, which trains the model to output a desc age order. The unbalance training data may lead the model to prefer outputting a desc order even its column is the 'date of birth'.

The second reason is that the model has insufficient generalization ability for similar problems. Many training examples belong to the T3 and T4. However, these examples can not cover all cases. For example, the training data may not or rarely contain examples where the USA is substituted with the United States, but we expect the models can still handle these examples correctly. The third reason is that a word will be used twice to generate schema items and SQL structure as we discussed the T5 in Section 2.2.

## 4 Conclusion

We introduce Spider-DK, a human-curated dataset based on the Spider benchmark for evaluating the generalization of text-to-SQL models, with the focus of understanding the domain knowledge. We demonstrate that the performance of existing text-to-SQL models drops dramatically on Spider-DK, even if the domain knowledge appears in the training set. Our evaluation indicates that the models do not always understand the underlying domain knowledge for prediction, thus we consider improving the model understanding of domain knowledge as an important direction to achieve the cross-domain generalization.

## Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments. Matthew Purver is partially supported by the EPSRC under grant EP/S033564/1, and by the European Union’s Horizon 2020 programme under grant agreement 825153 (EMBEDDIA, Cross-Lingual Embeddings for Less-Represented Languages in European News Media). Xinyun Chen is supported by the Facebook Fellowship. The results of this publication reflect only the authors’ views and the Commission is not responsible for any use that may be made of the information it contains.

## References

- Ben Bogin, Jonathan Berant, and Matt Gardner. 2019. [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.
- Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun, and Matthew Richardson. 2021. [Structure-grounded pretraining for text-to-SQL](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1337–1350, Online. 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yujian Gan, Xinyun Chen, Qiuping Huang, Matthew Purver, John R. Woodward, Jinxia Xie, and Pengsheng Huang. 2021. [Towards robustness of text-to-SQL models against synonym substitution](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2505–2515, Online. Association for Computational Linguistics.
- Yujian Gan, Matthew Purver, and John R. Woodward. 2020. [A review of cross-domain text-to-SQL models](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing: Student Research Workshop*, pages 108–115, Suzhou, China. Association for Computational Linguistics.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. 2019. [Towards Complex Text-to-SQL in Cross-Domain Database with Intermediate Representation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4524–4535, Florence, Italy. Association for Computational Linguistics.
- Binyuan Hui, Xiang Shi, Ruiying Geng, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu. 2021. [Improving text-to-sql with schema dependency learning](#). *CoRR*, abs/2103.04399.
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, Jayant Krishnamurthy, and Luke Zettlemoyer. 2017. [Learning a Neural Semantic Parser from User Feedback](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 963–973.
- Daniel Keysers, Nathanael Schärli, Nathan Scales, Hylke Buisman, Daniel Furrer, Sergii Kashubin, Nikola Momchev, Danila Sinopalnikov, Lukasz Stafiniak, Tibor Tihon, Dmitry Tsarkov, Xiao Wang, Marc van Zee, and Olivier Bousquet. 2020. [Measuring compositional generalization: A comprehensive method on realistic data](#). In *International Conference on Learning Representations*.
- Inbar Oren, Jonathan Herzig, Nitish Gupta, Matt Gardner, and Jonathan Berant. 2020. [Improving compositional generalization in semantic parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2482–2495, Online. Association for Computational Linguistics.
- 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)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Ohad Rubin and Jonathan Berant. 2021. [SmBoP: Semi-autoregressive bottom-up semantic parsing](#). In

- Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 311–324, Online. Association for Computational Linguistics.
- Peter Shaw, Ming-Wei Chang, Panupong Pasupat, and Kristina Toutanova. 2021. [Compositional generalization and natural language variation: Can a semantic parsing approach handle both?](#) In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 922–938, Online. Association for Computational Linguistics.
- Peng Shi, Patrick Ng, Zhiguo Wang, Henghui Zhu, Alexander Hanbo Li, Jun Wang, Cícero Nogueira dos Santos, and Bing Xiang. 2020. [Learning contextual representations for semantic parsing with generation-augmented pre-training.](#) *CoRR*, abs/2012.10309.
- Robyn Speer and Catherine Havasi. 2012. [Representing General Relational Knowledge in ConceptNet 5.](#) In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 3679–3686, Istanbul, Turkey. European Language Resources Association (ELRA).
- Alane Suhr, Ming-Wei Chang, Peter Shaw, and Kenton Lee. 2020. [Exploring unexplored generalization challenges for cross-database semantic parsing.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8372–8388, Online. Association for Computational Linguistics.
- Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 2020. [RAT-SQL: Relation-Aware Schema Encoding and Linking for Text-to-SQL Parsers.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7567–7578, Online. Association for Computational Linguistics.
- 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. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task.](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- John M Zelle and Raymond J Mooney. 1996. [Learning to Parse Database Queries Using Inductive Logic Programming.](#) In *Proceedings of the Thirteenth National Conference on Artificial Intelligence - Volume 2*, pages 1050–1055.
- Liang Zhao, Hexin Cao, and Yunsong Zhao. 2021. [GP: context-free grammar pre-training for text-to-sql parsers.](#) *CoRR*, abs/2101.09901.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. [Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning.](#) *CoRR*, abs/1709.0.