

Correcting Semantic Parses with Natural Language through Dynamic Schema Encoding

Parker Glenn, Parag Pravin Dakle, Preethi Raghavan

Fidelity Investments, AI Center of Excellence

{parker.glenn, paragpravin.dakle, preethi.raghavan}@fmr.com

Abstract

In addressing the task of converting natural language to SQL queries, there are several semantic and syntactic challenges. It becomes increasingly important to understand and remedy the points of failure as the performance of semantic parsing systems improve. We explore semantic parse correction with natural language feedback, proposing a new solution built on the success of autoregressive decoders in text-to-SQL tasks. By separating the semantic and syntactic difficulties of the task, we show that the accuracy of text-to-SQL parsers can be boosted by up to 26% with only one turn of correction with natural language. Additionally, we show that a T5-base model is capable of correcting the errors of a T5-large model in a zero-shot, cross-parser setting.

1 Introduction

The task of parsing natural language into structured database queries has been a long-standing benchmark in the field of semantic parsing. Success at this task allows individuals without expertise in the downstream query language to retrieve information with ease. This helps to improve data literacy, democratizing accessibility to otherwise opaque public database systems.

Many forms of semantic parsing datasets exist, such as parsing natural language to programming languages (Ling et al., 2016; Oda et al., 2015; Quirk et al., 2015), Prolog assertions for exploring a database of geographical data (Zelle and Mooney, 1996), or SPARQL queries for querying a large knowledge base (Talmor and Berant, 2018). The current work discusses parsing natural language into a structured query language (SQL), perhaps the most well-studied sub-field of semantic parsing.

Most text-to-SQL works frame the task as a one-shot mapping problem. Methods include transition-based parsers (Yin and Neubig, 2018), grammar-based decoding (Guo et al., 2019; Lin et al., 2019),




Figure 1: Example item from the SPLASH dataset. An incorrect parse from a neural text-to-SQL model is paired together with natural language feedback commenting on how the parse should be corrected.

and the most popular approach as of late, sequence to sequence (seq2seq) models (Scholak et al., 2021; Qi et al., 2022; Xie et al., 2022).

In contrast to the one-shot approach, conversational text-to-SQL aims to interpret the natural language to structured representations in the context of a multi-turn dialogue (Yu et al., 2019a,b). It requires some form of state tracking in addition to semantic parsing to handle conversational phenomena like coreference and ellipsis (Rui Zhang, 2019; Hui et al., 2021; Cai et al., 2022).

Interactive semantic parsing frames the task as a multi-turn interaction, but with a different objective than pure conversational text-to-SQL. As a majority of parsing mistakes that neural text-to-SQL parsers make are minor, it is often feasible for humans to suggest fixes for such mistakes using natural language feedback. Displayed in Figure 1, SPLASH (Semantic Parsing with Language

Assistance from Humans) is a text-to-SQL dataset containing erroneous parses from a neural text-to-SQL system alongside human feedback explaining how the interpretation should be corrected (Elgohary et al., 2020). Most similar to SPLASH is the INSPIRED dataset (Mo et al., 2022), which aims to correct errors in SPARQL parses from the ComplexWebQuestions dataset (Talmor and Berant, 2018). While the interactive semantic parsing task evaluates a system’s ability to incorporate human feedback, as noted in Elgohary et al. (2020), it targets a different modeling aspect than the traditional conversational paradigm. Hence, good performance on one does not guarantee good performance on the other task.

We make the following contributions: (1) We achieve a new state-of-the-art on the interactive parsing task SPLASH, beating the best published correction accuracy (Elgohary et al., 2021) by 12.33% using DestT5 (Dynamic Encoding of Schemas using T5); (2) We show new evidence that the decoupling of syntactic and semantic tasks improves text-to-SQL results (Li et al., 2023), proposing a novel architecture which leverages a single language model for both tasks; (3) We offer a new small-scale test set for interactive parsing¹, and show that a T5-base interactive model is capable of correcting errors made by a T5-large parser.

2 Dataset

In this work, we evaluate our models on the SPLASH dataset as introduced in Elgohary et al. (2020). It is based on Spider, a large multi-domain and cross-database dataset for text-to-SQL parsing (Yu et al., 2018). Incorrect SQL parses were selected from the output of a Seq2Struct model trained on Spider (Shin, 2019). Seq2Struct achieves an exact set match accuracy of 42.94% on the development set of Spider.

Alongside the incorrect parse, an explanation of the SQL query is generated using a rule-based template. Annotators were then shown the original question q alongside the explanation and asked to provide natural language feedback f such that the incorrect parse p' could be resolved to the final gold parse p .

Each item in the SPLASH dataset is associated with a relational database \mathcal{D} . Each database has a schema \mathcal{S} containing tables $T = \{t_1, t_2, \dots, t_N\}$ and columns $C = \{c_1^1, \dots, c_{n_1}^1, c_1^2, \dots, c_{n_2}^2, c_1^N, \dots, c_{n_N}^N\}$,

where N is the number of tables, and n_i is the number of columns in the i -th table. Figure 1 displays an example item from the SPLASH dataset, excluding the full database schema \mathcal{S} for brevity.

3 Model

3.1 Dynamic Schema Encoder

In converting natural language to SQL, a parser must handle both the semantic challenges in selecting the correct tables and columns from the database schema, and generate valid SQL syntax. As shown in Li et al. (2023), decoupling the schema linking and skeleton parsing tasks in text-to-SQL improves results when applied to the Spider dataset. We take a similar approach with the SPLASH dataset, separating the semantic and syntactic challenges of text-to-SQL by introducing an auxiliary schema prediction model. This auxiliary model serializes only the most relevant schema items into the input for the final seq2seq text-to-SQL model.

The task of the schema prediction is to output only those schema items (tables, columns, values) that appear in the gold SQL p . The inputs can be represented as follows.

$$d = t_1 : c_1^1, \dots, c_{n_1}^1 | \dots | t_N : c_1^N, \dots, c_{n_N}^N \quad (1)$$

$$x = ([CLS], q, [SEP], d, [SEP], p', [SEP], f) \quad (2)$$

Where d represents a flattened representation of the database schema \mathcal{S} , q is the question, p' is the incorrect parse from SPLASH, and f is the natural language feedback. For each schema item, the task is to predict the presence or absence of the item in the final gold SQL parse p .

By introducing this auxiliary schema prediction model, the final text-to-SQL model should only be tasked with stitching together the predicted schema items into valid SQL logic. As shown in the example in Figure 2, the text-to-SQL model is able to filter out the unnecessary “join” clauses from the incorrect parse, given the only table predicted by the schema prediction is “Flights”.

This approach was validated by carrying out a simple experiment. We serialize only those “gold” schema items that appear in the translated SQL and fine-tune a T5-base model² on the Spider dataset to achieve a best 78.10% execution accuracy. This

¹<https://github.com/parkervg/DestT5>

²<https://huggingface.co/tscholak/t5.1.1.lm100k.base>

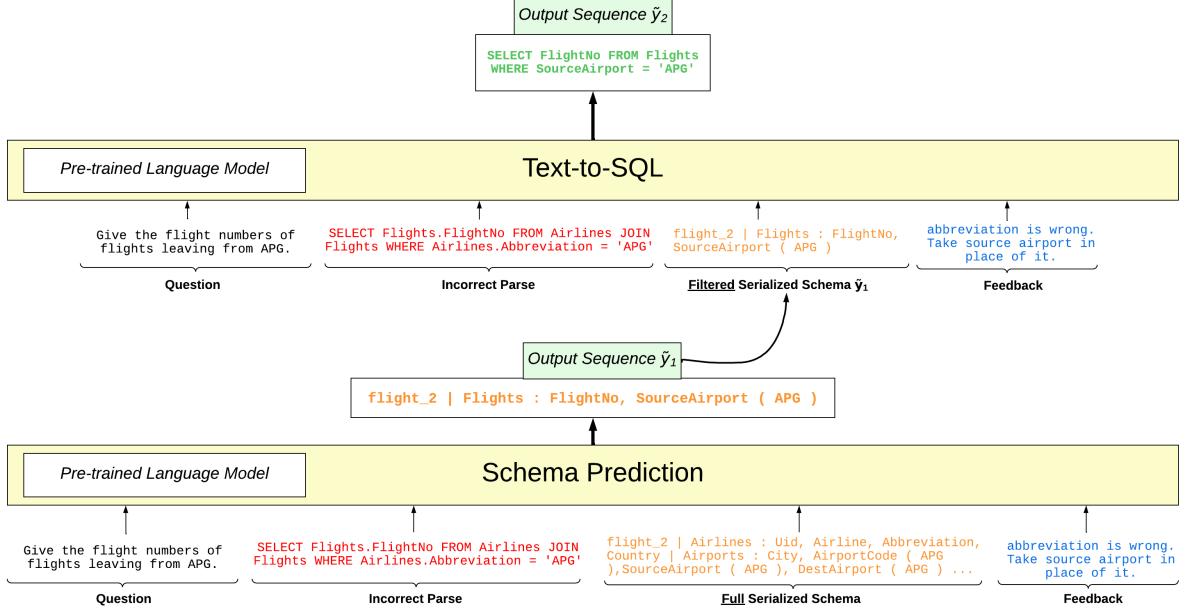


Figure 2: Model architecture. In “Schema Prediction”, the database schema is filtered to only the relevant items \tilde{y}_1 using a classifier or generator described in Section 3.1. In “Text-to-SQL”, the output of the schema prediction model is used to generate the final parse \tilde{y}_2 .

beats the vanilla T5-base model³ by 18.7%, demonstrating that successful schema prediction sets up a text-to-SQL model to predict the final query with high accuracy.

Schema Classifier We adopt the RoBERTa-large schema prediction described in Li et al. (2023) for our classification model. To alleviate the label imbalance problem caused by sparse schema targets, focal loss is used as the loss function (Lin et al., 2017). Focal loss adds a factor $(1 - p_t)^\gamma$ to standard cross entropy loss, reducing relative loss for well-classified examples and putting more focus on misclassified examples.

$$\mathcal{L}_2 = \frac{1}{N} \sum_{i=1}^N FL(y_i, \hat{y}_i) + \frac{1}{M} \sum_{i=1}^N \sum_{k=1}^{n_i} FL(y_k^i, \hat{y}_k^i) \quad (3)$$

Where FL denotes the focal loss function. y_i is the ground truth label of the i -th table, either 0 or 1 indicating the presence or absence, respectively. Similarly, y_k^i is the ground truth label of the k -th column in the i -th table.

Rather than using a hard probability threshold, hyperparameters k_1 and k_2 are introduced. Taking the probabilities from the cross-encoder, only

the top- k_1 tables and top- k_2 columns are kept and serialized into a ranked schema serialization, descending by probability.

Schema Generator In addition to the previously discussed RoBERTa-large cross-encoder, we also experiment with a generative schema prediction model. T5 (Text-to-Text Transfer Transformer) is a transformer-based encoder-decoder model that converts all NLP problems into a text-to-text format (Raffel et al., 2020). In our task setup, the encoder applies its bidirectional attention mechanism over the features from SPLASH and the serialized schema items, depicted in Equation 2. The decoder, then, generates the correct SQL parse, employing teacher forcing during the training phase. It is fine-tuned using standard cross-entropy loss.

$$L_1 = - \sum_{i=1}^M y_i \log(\hat{y}_i) \quad (4)$$

The target label y_i will always take the form of tokens comprising the gold schema items, i.e., those tables and columns that appear in the correct SQL parse. We format the multi-label targets y as text following the structure shown below. Note that this is the same structure we use to serialize the flattened database schema d in Equation 1.

[db_id] | [table] : [column] (...)

³<https://huggingface.co/tscholak/1zha5ono>

Schema Model	F1	Precision	Recall
Generator	88.98	90.84	89.18
Classifier	34.50	22.12	94.41

Table 1: Performance of schema prediction models in predicting gold schema items on the SPLASH test set. Note that the classification-based method of Li et al. (2023) trades low precision for high recall⁵.

As the theoretical output space of \hat{y} is the unconstrained vocabulary of the T5 model, schema hallucinations are possible, and column/table pairs may be generated that do not exist in the database context⁴. A trade-off in this approach, however, is that the generation objective allows us to bypass the need for hyperparameters k_1 and k_2 , as we simply keep the greedy argmax of \hat{y} directly at each timestep. As shown in Table 1, this optimization objective results in far greater precision than the classification approach but suffers a drop in recall.

3.2 Text-to-SQL Encoder/Decoder

We use a T5-base model to encode the unified input (with schema predictions) and generate the SQL query (Raffel et al., 2020).

3.3 SQL Normalization

We follow the same normalization procedure described in Li et al. (2023). Specifically, we normalize both the incorrect parses and gold SQL queries by (1) replacing table aliases with their original names, (2) adding an *ASC* keyword if *ORDER BY* doesn’t already specify, (3) lower-casing all text, and (4) adding spaces around parentheses and replacing double quotes with single quotes.

4 Experiments

4.1 Experimental Setup

We run a series of experiments on the SPLASH dataset to evaluate the robustness of the proposed method. The training set contains 2,775 unique questions from the train split of Spider. SPLASH annotators were also asked to generate paraphrases for a single piece of feedback to improve diversity, resulting in a total of 7,481 items in the train split. The SPLASH test set is based on 506 items from

⁴We note that Scholak et al. (2021) offers a solution for these schema hallucinations, but leave the integration of Picard to future work.

⁵Not considered in this table is the ranking-enhanced nature of the RoBERTa-large method.

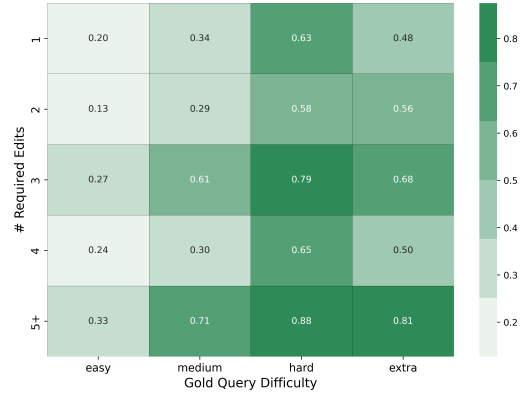


Figure 3: DestT5 error rates on the SPLASH test set, using the Spider exact match metric. As the distance (*# Required Edits*) from the incorrect parse to the gold query increases, error rates also increase.

the Spider dev split, coming out to 962 total test items with paraphrasing.

4.2 Evaluation Metrics

Exact Set Match (EM) This metric evaluates the structural correctness of the predicted SQL. It checks for an orderless set match between each component in the predicted and gold query, ignoring predicted values. Many early text-to-SQL models only report EM accuracy.

Execution Accuracy (EX) Execution accuracy compares the execution results of the predicted SQL query and the gold SQL query. Since two SQL queries that do not have an exact set match may execute to the same results (e.g. “...ORDER BY val ASC LIMIT 1” and “SELECT MAX(val)”), this metric serves as a performance upper bound. However, this metric can suffer from a high false positive rate. For this reason, we use the test suite execution accuracy with optimized database values described in Zhong et al. (2020).

4.3 Implementation Details

Text-to-SQL All text-to-SQL models use a fine-tuned T5-base. We use the same hyperparameters specified in the PICARD codebase⁶. Models were fine-tuned with Adafactor (Shazeer and Stern, 2018) with a learning rate 1e-4, batch size 16 for 256 epochs. A linear warm-up for the first 10% of training steps is employed, followed by cosine decay.

⁶<https://github.com/ServiceNow/picard>

			Shuffled Feature EM% Change	
	Schema Model	EM%	Feedback	Incorrect Parse
All	None	41.17	-	-
	Generator	51.35	-2.17	-28.27
	Classifier	49.79	-2.7	-11.64
- Question	Generator	48.96	-4.47	-30.77
	Classifier	35.97	-11.23	-29.94
- Explanation	Generator	53.43	-1.77	-18.09
	Classifier	49.27	-2.08	-17.57
- Question	Generator	47.00	-5.53	-38.68
- Explanation	Classifier	38.98	-12.47	-36.9

Table 2: Results on SPLASH test set with various features and schema prediction models. *Generator* refers to the T5-large model, and *Classifier* refers to the RoBERTa-large model of Li et al. (2023). The models are evaluated on the test set with shuffled features to examine the extent to which they utilize the unique interactive components of the parsing task. In bold is DestT5.

Schema Generation T5-large was used for the schema generation model. It was fine-tuned using Adafactor with a constant learning rate of 1e-4 and a batch size of 4 for 512 epochs.

Schema Classification For the schema classification model, we follow the implementation and hyperparameters described in Li et al. (2023). Specifically, we train a cross-encoder based on RoBERTa-large (Liu et al., 2019). AdamW (Loshchilov and Hutter, 2017) with a batch size of 32 and a learning rate of 1e-5 is used for optimization. Focal loss is used to alleviate the label-imbalance problem that comes from sparse schema targets. The threshold hyperparameters k_1 and k_2 are set to 4 and 5, respectively. Specifically, only the top-4 tables and top-5 columns with the highest logits are kept and serialized as a ranked input to the text-to-SQL model.

4.4 Evaluation

Unlike the Spider dataset, performance on the SPLASH dataset is more nuanced and must be viewed holistically. To this end, we plot both “Exact Match %” and “Shuffled Feature Change” in Table 2. The ideal model is one that achieves a competitive exact match metric, while experiencing a large drop in performance with shuffled feedback and incorrect parses⁷. We find the highest exact match accuracy when removing the explanation of the incorrect parse, and by using a T5-based

⁷We note that a T5-base model fine-tuned with the Spider train set achieves 50.00 EM on the SPLASH test set.

generative schema prediction model. This model, denoted in bold in Table 2, is later referred to as DestT5 (Dynamic Encoding of Schemas using T5). Achieving an EM score of **53.43%**, DestT5 beats the previous best score of NL-EDIT by 12.33% (Elgohary et al., 2021).

Using the scripts provided from Elgohary et al. (2021) to count SQL edits, we plot error rates on the SPLASH test set for both gold query difficulty and the number of edits. “Difficulty” is defined by Yu et al. (2018) and classifies each SQL query into one of four categories depending on the complexity of the query. As seen in the heatmap, error rates share a positive correlation with both SQL difficulty and # edits required to reach the gold parse.

4.5 Generalizing to Other Parsers

In recent years, massive strides have been made in the task of semantic parsing. Since the release of the SPLASH dataset, variations of T5 have largely taken the top spots in the Spider leaderboard. As of April 2023, all 6 models in the top 10 with corresponding publications build off of some T5 model. It is fair, then, to ask if performance on the SPLASH dataset actually corresponds to the ability to fix errors made with modern parsing systems, such as those utilizing T5.

To this end, we evaluate DestT5 on the crowd-sourced test sets⁹ based on errors made by EditSQL (Rui Zhang, 2019), TaBERT (Yin et al., 2020), and RAT-SQL (Wang et al., 2020). Additionally, we

⁹<https://github.com/MSR-LIT/NLEdit>

	Seq2Struct (SPLASH)	EditSQL	TaBERT	RAT-SQL	T5-Large
Spider Dev EM%	41.3	57.6	65.2	69.7	71.2
Spider Dev EX%	-	-	-	-	74.4
NL-EDIT					
SPLASH Test Set EM%	41.1	28	22.7	21.3	-
SPLASH Test Set EX%	-	-	-	-	-
EM Δ w/ Interaction	+20.3	+8.9	+5.9	+4.3	-
EX Δ w/ Interaction	-	-	-	-	-
DESTT5 (OURS)					
SPLASH Test Set EM%	53.43	31.82	31.47	28.37	26.1
SPLASH Test Set EX%	56.86	40.3	28.84	36.53	30.43
EM Δ w/ Interaction	+26.15	+10.16	+8.13	+5.71	+2.83
EX Δ w/ Interaction	-	-	-	-	+3.3

Table 3: Evaluating zero-shot generalization of DestT5 to other modern parsers. Shown are the scores without interaction on the full Spider dev set, as well as the Δ w/ **Interaction** on the Spider dev set following single-turn corrections with NL-EDIT and DESTT5. This change is a byproduct of the size of the test sets (962, 330, 267, 208, and 112 left-to-right), and it is expected to increase proportional to the reported **Test Set EM%/EX%** as the size of the dataset increases. We indicate instances where the scores are not publicly available for a given model with -.

Text-to-SQL Model	Schema F1	# Hallucinated Schema Items
T5-large ⁸	79.00	92
T5-base	73.92	121
DestT5	80.09	59

Table 4: Analysis of the schema items produced by the final text-to-SQL model. DestT5, with an auxiliary schema prediction model, identifies the presence of gold schema items with a higher F1 than a T5-large text-to-SQL model alone.

compile a new, small-scale test set of errors made by a fine-tuned T5-large model¹⁰ on the Spider dev set. It contains 112 items annotated with feedback referencing the erroneous parse made by the model and is later referred to as the “T5-large Test Set”.

Table 3 plots the end-to-end accuracy of DestT5. As mentioned in Elgohary et al. (2021), there is a notable drop in the end-to-end gains as the accuracy of the base parser improves. This is likely due to the fact that as parsers improve, most of the errors are based on very complex gold SQL queries.

4.6 Error Analysis

4.7 Errors on T5-Large Test Set

Figure 4 depicts the outputs of a randomly selected set of interactions from the T5-large test set. We discuss some of the examples below.

¹⁰<https://huggingface.co/tscholak/3vnuv1vf>

In Example 1, the original T5-large text-to-SQL model fails to map the phrase “all lines” to both columns *line_1* and *line_2*. However, even with the feedback “Find *line_2* as well”, the auxiliary schema prediction model fails to select “*line_2*” as a schema candidate. As a result, the final DestT5 text-to-SQL is not equipped with enough context to generate the correct parse.

In Example 2, an ‘easy’ gold query (“SELECT MIN(loser_rank) FROM matches”) is incorrectly parsed. This is likely due to the same reason described in Lin et al. (2020), characterized by difficulty in mapping “predominantly” to *spoken by the largest percentage of the population*: it remains challenging for large pre-trained models to ground terms like “best rank” to the DB schema. Pre-training tasks have been proposed in attempts to further improve schema grounding in LLMs, but more work can be done to align LLMs with lexical

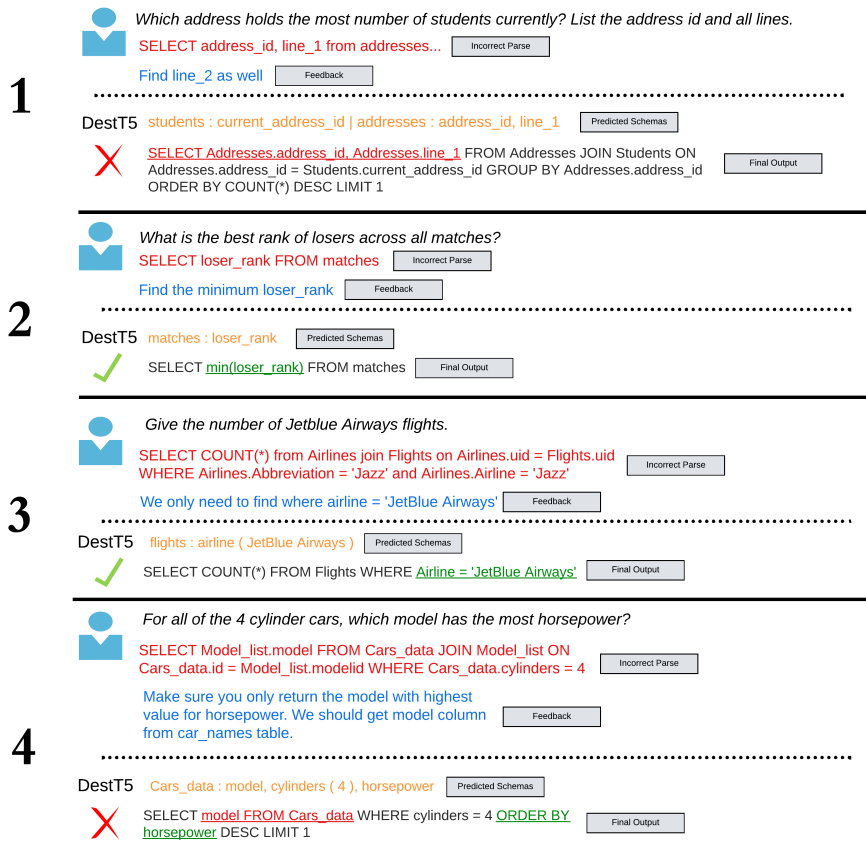


Figure 4: Example outputs of DestT5 on errors made with a T5-large text-to-SQL model. When the schema prediction model fails to identify schema items, the final text-to-SQL output is incorrect. However, when the schema prediction model is correct, it allows the text-to-SQL component to focus its efforts on generating valid SQL syntax, faithful to the feedback. See section 4.7 for more detailed analysis of these examples.

constructs grounded to the syntax of semantic parsing tasks (Deng et al., 2021; Yin et al., 2020). In one turn of interaction with DestT5, this syntax error is corrected.

Example 4 displays an interaction parsing long feedback with mixed success. The interaction allows DestT5 to remedy the missed semantic mapping from “most horsepower” to the “ORDER BY horsepower” clause, but it hallucinates the “Cars_data” from the “model” table, failing to learn from the feedback saying otherwise.

5 Discussion

5.1 Impact of Auxiliary Schema Prediction

Table 2 displays the EM of a standard text-to-SQL model with no auxiliary schema prediction (with all schema items directly serialized as input). As shown, the score drops from 51.35% with an auxiliary generator to 41.17% without. We hypothesize that given the increased number of features in interactive semantic parsing (explanation, feedback,

incorrect parse), distilling the role of the text-to-SQL model to primarily handling syntax parsing prevents excessive proliferation of feature interactions.

Table 4 displays the schema F1 scores of various text-to-SQL models. Schema F1 is calculated by comparing those schema items (tables, columns) generated in the predicted parse to the schema items in the gold SQL. As shown, implementing a dedicated schema prediction model into a text-to-SQL pipeline helps identify those gold schema items with a higher F1 score, and minimizes schema hallucinations (i.e. generating tables/columns not present in the database schema).

How often does the text-to-SQL model use the predicted schemas?

We evaluate the usage rates of the predicted schema items by the final text-to-SQL model. Specifically, we examine the rate at which DestT5 either predicts a schema item not directly serialized by the schema prediction model, or fails to integrate a schema item that was serial-

ized. We find that on the SPLASH test set, there are 112 instances of overpredictions by the text-to-SQL model and 210 underpredictions. There is an average distance of 0.81 between the serialized schema items and gold schema items, and 0.93 between the schema items predicted by the text-to-SQL model and gold. This indicates that, if the text-to-SQL model were explicitly restricted to use only the schema items generated by the auxiliary schema prediction model, performance will improve. We leave this and other combinations of the two models (such as joint training) to future work.

5.2 Evaluating Interactive Parsing

The goal of interactive semantic parsing is not to parse the most interactions correctly on the SPLASH test set, but more specifically to parse those interactions correctly that the original text-to-SQL model parsed incorrectly. For example, if a hypothetical interactive parsing model A achieves a high EM% on the SPLASH test set, but the “ Δ w/ Interaction” metric with modern parsers is small, then the model serves minimal utility in an actual conversational setting. On the other hand, if a model B performs poorly on the SPLASH test set but demonstrates a high “ Δ w/ Interaction”, we would deem this model as the better interactive semantic parser.

We argue, then, that the “Correction Acc. (%)” metric from SPLASH should be replaced in favor of the end-to-end accuracy, referred to as “ Δ w/ Interaction” in Elgohary et al. (2021).

Specifically, future work should include Execution Accuracy (EX%) along with Exact Set Match (EM%). As the set of errors made by modern parsers increasingly drifts towards more difficult gold SQL parses, it becomes more likely that the EM% and EX% scores will be disjoint. Examining the errors by T5-large, it was common for a gold parse to be expressed with an “EXCEPT SELECT” clause, whereas the predicted SQL executed identically with a “NOT IN” clause.

Additionally, as depicted in Table 3, the EX% score is higher than EM% for all test sets except for TaBERT. This is due to the fact that TaBERT does not predict values. Instead, it uses the placeholder “value” instead of string values, and “LIMIT 0” in limit clauses¹¹. Though these instances are not

¹¹We find this odd, as the feedback provided in the TaBERT test set comments on the values

judged as incorrect with EM, they are penalized with EX.

6 Conclusion

We present a new model, DestT5 (Dynamic Encoding of Schemas using T5), which achieves a new state-of-the-art correction accuracy on the interactive parsing dataset SPLASH. By using T5 as a schema prediction model, we display better performance compared to classification-based methods. We validate our results on a new test set for interactive semantic parsing based on a modern parser, and offer recommendations for evaluating future systems.

Limitations

As mentioned in Table 3, one limitation of the current study is the small scale of the test sets with modern parsers. We encourage future work to emphasize the development and evaluation on these test sets, specifically those which more closely reflect the current SoTA in text-to-SQL (e.g. T5). Additionally, though we have shown using an auxiliary schema prediction model greatly improves the performance of a text-to-SQL system, the addition of a model for the text-to-SQL task is a limitation given the time and training resources required.

References

- Zefeng Cai, Xiangyu Li, Binyuan Hui, Min Yang, Bowen Li, Binhua Li, Zheng Cao, Weijie Li, Fei Huang, Luo Si, and Yongbin Li. 2022. [STAR: SQL guided pre-training for context-dependent text-to-SQL parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1235–1247, Abu Dhabi, United Arab Emirates. 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.
- Ahmed Elgohary, Saghar Hosseini, and Ahmed Hassan Awadallah. 2020. [Speak to your Parser: Interactive Text-to-SQL with Natural Language Feedback](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2065–2077, Online. Association for Computational Linguistics.

- Ahmed Elgohary, Christopher Meek, Matthew Richardson, Adam Fourney, Gonzalo Ramos, and Ahmed Hassan Awadallah. 2021. [NL-EDIT: Correcting semantic parse errors through natural language interaction](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5599–5610, Online. Association for Computational Linguistics.
- Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jianguang Lou, Ting Liu, and Dongmei Zhang. 2019. Towards complex text-to-sql in cross-domain database with intermediate representation. In *Proceeding of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics.
- Binyuan Hui, Ruiying Geng, Qiyu Ren, Binhua Li, Yongbin Li, Jian Sun, Fei Huang, Luo Si, Pengfei Zhu, and Xiaodan Zhu. 2021. Dynamic hybrid relation exploration network for cross-domain context-dependent semantic parsing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13116–13124.
- Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. 2023. Resdsq: Decoupling schema linking and skeleton parsing for text-to-sql. In *AAAI*.
- Kevin Lin, Ben Bogin, Mark Neumann, Jonathan Berant, and Matt Gardner. 2019. Grammar-based neural text-to-sql generation. *arXiv preprint arXiv:1905.13326*.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988.
- Xi Victoria Lin, Richard Socher, and Caiming Xiong. 2020. [Bridging Textual and Tabular Data for Cross-Domain Text-to-SQL Semantic Parsing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4870–4888, Online. Association for Computational Linguistics.
- Wang Ling, Phil Blunsom, Edward Grefenstette, Karl Moritz Hermann, Tomáš Kočiský, Fumin Wang, and Andrew Senior. 2016. [Latent predictor networks for code generation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 599–609, Berlin, Germany. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Lingbo Mo, Ashley Lewis, Huan Sun, and Michael White. 2022. [Towards transparent interactive semantic parsing via step-by-step correction](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 322–342, Dublin, Ireland. Association for Computational Linguistics.
- Yusuke Oda, Hiroyuki Fudaba, Graham Neubig, Hideaki Hata, Sakriani Sakti, Tomoki Toda, and Satoshi Nakamura. 2015. Learning to generate pseudo-code from source code using statistical machine translation. In *2015 30th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 574–584. IEEE.
- Jiexing Qi, Jingyao Tang, Ziwei He, Xiangpeng Wan, Yu Cheng, Chenghu Zhou, Xinbing Wang, Quanshi Zhang, and Zhouhan Lin. 2022. [RASAT: Integrating Relational Structures into Pretrained Seq2Seq Model for Text-to-SQL](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3215–3229, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chris Quirk, Raymond Mooney, and Michel Galley. 2015. [Language to code: Learning semantic parsers for if-this-then-that recipes](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 878–888, Beijing, China. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- He Yang Er Sungrok Shim Eric Xue Xi Victoria Lin Tianze Shi Caiming Xiong Richard Socher Dragomir Radev Rui Zhang, Tao Yu. 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*, Hong Kong, China.
- Torsten Scholak, Nathan Schucher, and Dzmitry Bahdanau. 2021. [PICARD: Parsing incrementally for constrained auto-regressive decoding from language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9895–9901, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Noam Shazeer and Mitchell Stern. 2018. Adafactor: Adaptive learning rates with sublinear memory cost. In *International Conference on Machine Learning*, pages 4596–4604. PMLR.
- Richard Shin. 2019. [Encoding Database Schemas with Relation-Aware Self-Attention for Text-to-SQL Parsers](#). ArXiv:1906.11790 [cs, stat].

- Alon Talmor and Jonathan Berant. 2018. [The Web as a Knowledge-Base for Answering Complex Questions](#). 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 641–651, New Orleans, Louisiana. 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.
- Tianbao Xie, Chen Henry Wu, Peng Shi, Ruiqi Zhong, Torsten Scholak, Michihiro Yasunaga, Chien-Sheng Wu, Ming Zhong, Pengcheng Yin, Sida I. Wang, Victor Zhong, Bailin Wang, Chengzu Li, Connor Boyle, Ansong Ni, Ziyu Yao, Dragomir Radev, Caiming Xiong, Lingpeng Kong, Rui Zhang, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2022. [UnifiedSKG: Unifying and multi-tasking structured knowledge grounding with text-to-text language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 602–631, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Pengcheng Yin and Graham Neubig. 2018. [TRANX: A transition-based neural abstract syntax parser for semantic parsing and code generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 7–12, Brussels, Belgium. Association for Computational Linguistics.
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. [TaBERT: Pretraining for joint understanding of textual and tabular data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8413–8426. Online. Association for Computational Linguistics.
- Tao Yu, Rui Zhang, Heyang Er, Suyi Li, Eric Xue, Bo Pang, Xi Victoria Lin, Yi Chern Tan, Tianze Shi, Zihan Li, Youxuan Jiang, Michihiro Yasunaga, Sungrok Shim, Tao Chen, Alexander Fabri, Zifan Li, Luyao Chen, Yuwen Zhang, Shreya Dixit, Vincent Zhang, Caiming Xiong, Richard Socher, Walter Lasecki, and Dragomir Radev. 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, Hong Kong, China. 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.
- Tao Yu, Rui Zhang, Michihiro Yasunaga, Yi Chern Tan, Xi Victoria Lin, Suyi Li, Heyang Er, Irene Li, Bo Pang, Tao Chen, Emily Ji, Shreya Dixit, David Proctor, Sungrok Shim, Jonathan Kraft, Vincent Zhang, Caiming Xiong, Richard Socher, and Dragomir Radev. 2019b. [SPaC: Cross-domain semantic parsing in context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4511–4523, Florence, Italy. Association for Computational Linguistics.
- John M Zelle and Raymond J Mooney. 1996. Learning to parse database queries using inductive logic programming. In *Proceedings of the national conference on artificial intelligence*, pages 1050–1055.
- Ruiqi Zhong, Tao Yu, and Dan Klein. 2020. [Semantic evaluation for text-to-SQL with distilled test suites](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 396–411. Online. Association for Computational Linguistics.