

Explaining Large Language Model-Based Neural Semantic Parsers (Student Abstract)

Daking Rai¹, Yilun Zhou², Bailin Wang², Ziyu Yao¹

¹ George Mason University,

² Massachusetts Institute of Technology

drai2@gmu.edu, yilun@mit.edu, bailinw@mit.edu, ziyuyao@gmu.edu

Abstract

While large language models (LLMs) have demonstrated strong capability in structured prediction tasks such as semantic parsing, few amounts of research have explored the underlying mechanisms of their success. Our work studies different methods for explaining an LLM-based semantic parser and qualitatively discusses the explained model behaviors, hoping to inspire future research toward better understanding them.

Introduction

Semantic parsing is a task of mapping natural language utterances to their logical forms like SQL queries or lambda expressions for database or knowledge base querying. Despite its structured prediction nature, recent work has shown that a large language model (LLM) which generates output sequentially could achieve comparable or even better performance than the traditional structured decoders (Scholak, Schucher, and Bahdanau 2021). However, why these LLMs could do well in semantic parsing is still unclear.

In this paper, we seek to provide one of the first studies toward explaining LLM-based neural semantic parsers. We use the text-to-SQL semantic parsing task (Yu et al. 2018) and the UnifiedSKG model (Xie et al. 2022) for a case study. We empirically explore a set of local explanation methods and quantitatively discussed the explanation results.

Method

(1) LIME (Ribeiro, Singh, and Guestrin 2016) generates an explanation by training locally-faithful interpretable models with the dataset obtained by perturbing the prediction instance. (2) Shapley value measures the importance of a feature by its average marginal contribution to the prediction score (Shapley 1952). (3) Kernel SHAP (Lundberg and Lee 2017) is another efficient way of estimating Shapley values by training a linear classifier. (4) LERG (Tuan et al. 2021) is a set of two approaches, LERG.L and LERG.S, recently adapted from LIME and Shapley value to conditioned sequence generation tasks. When applying these methods to explain an LLM-based semantic parser, we consider each output token as one prediction and attribute it to the input

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Necessity (left) and Sufficiency (right) scores when removing or keeping the top-K% important features.

features. (5) Attention: Prior work has revealed that attention may be interpreted as feature importance. Therefore, we also introduce an attention-based local explanation method, where the feature attribution is calculated by averaging the last layer of the multi-headed cross-attention weights.

Experimental Setup

In our experiments, we consider the task of text-to-SQL semantic parsing where the goal is to generate a SQL query given a natural language question and the database schema (i.e., tables and columns included in the database) as input. We experiment with UnifiedSKG (Xie et al. 2022), one of the state-of-the-art models, which adopts a T5 encoder-decoder structure.¹ Following Xie et al. (2022), we train and evaluate the parser on the Spider dataset (Yu et al. 2018).

Through the experiments, we seek to answer two *Research Questions (RQs)*: (1) *Which local explanation method is the most faithful to explaining the LLM-based UnifiedSKG parser?* (2) *How well does the explanation align with human intuitions?* To answer RQ1, we follow Tuan et al. (2021) and compare different explanation methods on two metrics: (a) Sufficiency measures the perplexity when keeping only the top-K% most important features by each explanation method; the lower the better/fairful. (b) Necessity measures the perplexity change when the top-K%

¹We used the “T5_base_prefix_spider_with_cell_value” version from <https://github.com/HKUNLP/UnifiedSKG>. We did not use the T5-3B version because of the large computational demand, which we will discuss in Section .

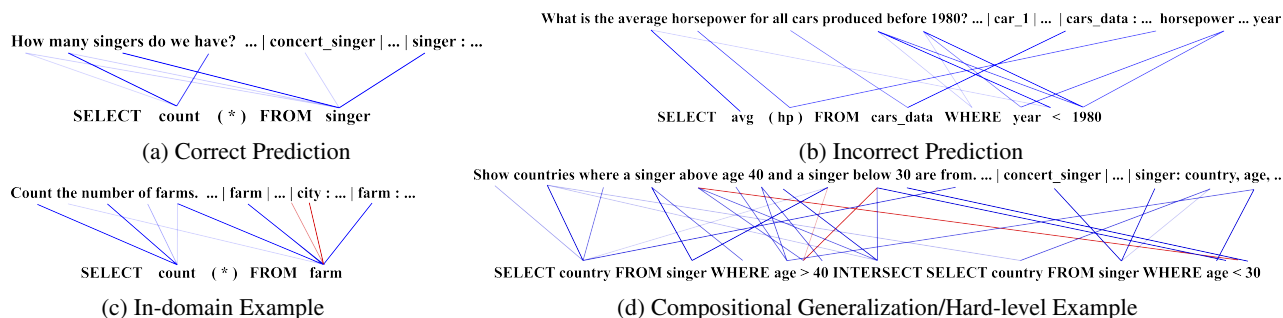


Figure 2: UnifiedSKG generally shows plausible explanations. In (a), “concert_singer” is the database name; “singer : ...” shows the table along with its columns (omitted with ellipsis); similarly for other examples. For (b), the selected column “hp” is incorrect and should be “horsepower”. Blue/red indicates positive/negative importance. Darkness indicates strength.

most important features are removed; the higher the better-/faithful. To answer RQ2, we qualitatively discuss the most faithful explanation results.

Experimental Results

Faithfulness. The results in Figure 1 show LERG_S has the best performance as per both sufficiency and necessity metrics with Kernel SHAP having comparable performance as well. In general, we observe that Shapley value-based explanation methods have more faithful explanations than other methods. In addition, we also found that attention-based explanations are more faithful than LIME and LERG_L.

Plausibility. Using LERG_S as a lens, we qualitatively study how UnifiedSKG works. We define *plausible* explanations as those which align well with human intuition. In our study, we classify each explanation into plausible or partially plausible ones. Interestingly, we didn’t find any explanation that is completely implausible. Under this setup, we investigate the four aspects listed below (Figure 2): **(1) Feature Attribution for (In)correct Predictions:** We randomly sample 20 examples where the model makes correct and incorrect predictions, respectively. We find out that in most cases (85% for correct and 70% for incorrect), LERG_S generates a plausible explanation for both types. **(2) Different Hardness Levels:** We randomly sample 20 examples for each hardness level – easy, medium, hard, and extra hard, as defined by the Spider benchmark based on the SQL complexity. We observed that in most cases the model behaviors are in line with human intuitions even at the extra hard level (80%; >90% for other levels). **(3) Compositional Generalization:** We seek to understand whether the model attributes the output fragments to correct features compositionally when it makes correct predictions. We conducted a similar manual examination as before and observed that in most (80%) cases our model shows compositionally generalizable feature attribution. **(4) In-domain vs. Out-of-domain:** As the Spider training and dev sets are split by databases (which could be seen as different domains), we also manually compare the model explanations in in-domain and out-of-domain cases. We observe that for both cases (75% and 80% respectively), the generated explanations were mostly plausible.

Discussion and Future Directions

Our study has revealed several challenges and opportunities in explaining an LLM-based semantic parser: **(1) Computational costs:** Most local explanation methods require model inference over a large set of input perturbations, which is computationally inefficient. Future work may look into improving the attention-based explanation method, which does not rely on perturbations and hence could save much computation. **(2) Feature interaction:** Traditional feature attribution does not provide information about how features (e.g., question tokens and contextual database schema items) interact with each other. Future work may uncover these interactions to gain deeper insights into how the model works. **(3) Explanation for user understanding:** Current saliency maps encompass a lot of information. Future work could examine how to present the information in a concise and friendly way such that users could easily grasp the intuition of the model prediction and verify its correctness. **(4) Explanation for debugging:** Future work should also investigate how the local explanation results could be used to probe and debug a semantic parser, such as to improve their capability in compositional generalization.

References

- Lundberg, S.M.; and Lee, S.-I.. 2017. A unified approach to interpreting model predictions. In *NeurIPS*.
- Ribeiro, M.T.; Singh, S.; Guestrin, C.. 2016. “Why should I trust you?”: Explaining the predictions of any classifier. In *SIGKDD*.
- Scholak T.; Schucher N.; Bahdanau D.. 2021. PICARD: Parsing incrementally for constrained auto-regressive decoding from language models. In *EMNLP*.
- Shapley, L.S.. 1952. A Value for N-Person Games. In *RAND Corporation*.
- Tuan, Y.; Pryor, C.; Chen, W.; et al. 2021. Local explanation of dialogue response generation. In *NeurIPS*.
- Xie, T; Wu, C.H.; Shi, P.; et al. 2022. UnifiedSKG: Unifying and multi-tasking structured knowledge grounding with text-to-text language models. In *EMNLP*.
- Yu, T.; Zhang, R.; Yang, K.; et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task. In *EMNLP*.