



Paraphrasing Techniques for Maritime QA system

Author

Shiri, F, Zhuo, TY, Li, Z, Pan, S, Wang, W, Haffari, R, Li, YF, Nguyen, V

Published

2022

Conference Title

2022 25th International Conference on Information Fusion (FUSION)

Version

Accepted Manuscript (AM)

DOI

[10.23919/FUSION49751.2022.9841287](https://doi.org/10.23919/FUSION49751.2022.9841287)

Rights statement

© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Downloaded from

<http://hdl.handle.net/10072/421077>

Griffith Research Online

<https://research-repository.griffith.edu.au>

Paraphrasing Techniques for Maritime QA system

Fatemeh Shiri[§], Terry Yue Zhuo[§], Zhuang Li[§],
Shirui Pan, Weiqing Wang, Reza Haffari, Yuan-Fang Li
Department of Data Science and AI
Faculty of Information Technology
Monash University, Melbourne, Australia
firstname.lastname@monash.edu

Van Nguyen
Defence Science and Technology Group
Adelaide, Australia
Van.Nguyen5@dst.defence.gov.au

Abstract—There has been an increasing interest in incorporating Artificial Intelligence (AI) into Defence and military systems to complement and augment human intelligence and capabilities. However, much work still needs to be done toward achieving an effective human-machine partnership. This work is aimed at enhancing human-machine communications by developing a capability for automatically translating human natural language into a machine-understandable language (e.g., SQL queries). Techniques toward achieving this goal typically involve building a *semantic parser* trained on a very large amount of high-quality manually-annotated data. However, in many real-world Defence scenarios, it is not feasible to obtain such a large amount of training data. To the best of our knowledge, there are few works trying to explore the possibility of training a semantic parser with limited manually-paraphrased data, in other words, zero-shot. In this paper, we investigate how to exploit paraphrasing methods for the automated generation of large-scale training datasets (in the form of paraphrased utterances and their corresponding logical forms in SQL format) and present our experimental results using real-world data in the maritime domain.

I. INTRODUCTION

In recent years, there has been an increasing interest in incorporating Artificial Intelligence (AI) into Defence and military systems to complement and augment human intelligence and capabilities. For instance, AI technologies for Intelligent, Surveillance and Reconnaissance (ISR) now play a significant role in maintaining situational awareness and assist human partners with decision making. While AI technologies are outstanding at handling the enormous volumes of data, pattern recognition and anomaly detection, humans excel at making decisions on limited data and sense when data has been compromised (as witnessed with adversarial machine learning). A collaboration between machines and humans therefore holds great promise to help increasing adaptability, flexibility and performance across many Defence contexts. However much work still needs to be done toward achieving an effective human-machine partnership, including addressing the bottlenecks of communication, comprehension and trust. In this work, we are focused on the communication aspect which is at the center of human-machine coordination and success.

Traditional methods require humans to interact with a machine using a machine language or controlled natural language (CNL) (e.g., [30]). However, there is often no time for personnel to translate input to a machine and interpret complex machine outputs in military operations. In our previous

work, we proposed Maritime DeepDive [34], an automated construction of knowledge graphs from unstructured natural language data sources as part of the RUSH project [22] for situational awareness. We are interested in endowing our situational awareness system with a capability to assist human decision makers by allowing them to interact with the system using human natural language.

In this work, we investigate the application of state-of-the-art techniques in AI and natural language processing (NLP) for the automated translation of human natural language questions into SQL queries, toward enhancing human-machine partnership. Specifically, we build a semantic parser in the maritime domain starting with zero manually annotated training examples. Moreover, with the help of paraphrasing techniques, we reduce the gap between synthesized utterances and the natural real-world utterances. Our contributions are,

- An efficient and effective approach for generating large-scale domain specific training dataset for Question Answering (QA) systems. Our training dataset includes synthetic samples and paraphrased samples obtained from state-of-the-art (SOTA) techniques. Therefore, the parser trained on such samples does not suffer from fundamental mismatches between the distributions of the automatically generated examples and the natural ones issued by real users.
- Evaluation and analyses of various paraphrasing techniques to produce diverse natural language questions.
- Discussion of lessons learned and recommendations.

II. BACKGROUND

A. Semantic Parsing

Semantic parsing is a task of translating natural language utterances into formal meaning representations, such as SQL and abstract meaning representations (AMR) [1]. Modern neural network approaches usually formulate semantic parsing as a machine translation problem and model the semantic parsing process with variations of Seq2Seq [35] frameworks. The output of the Seq2Seq models are either linearized LF (Logic Form) token sequences [6], [7], [39] or action sequences that construct LFs [4], [9], [17], [18], [27], [36]. Semantic parsing has a wide range of applications including question answering [4], [9], [36], programming synthesis [27], and natural language understanding (NLU) in dialogue systems [10].

Our work is mainly focused on question answering. The user questions expressed in natural language are converted

[§]Equal contribution

into SQL queries which are then executed on our Maritime DeepDive Knowledge Graph [34] to retrieve answers to the questions. The recent surveys [13], [16], [43] cover comprehensive reviews about recent semantic parsing studies. Of particular relevance to this work are *Bootstrapping Semantic Parsers*.

B. Bootstrapping Semantic Parsers

Data scarcity is always a serious problem in semantic parsing since it is difficult for the annotators to acquire expert knowledge about the meaning of the target representations. To solve this problem, one line of research is to bootstrap semantic parsers with semi-automatic data synthesis methods. [37], [39] use a set of synchronous context-free grammar (SCFG) rules and canonical templates to generate a large number of *clunky* utterance-SQL pairs, respectively. And then they hire crowd-workers to paraphrase the *clunky* utterances into natural questions. In order to reduce the paraphrase cost, [11] applies a paraphrase detection to automatically align the clunky utterances with the user query logs. This method requires access to user query logs, which is infeasible in many scenarios. In addition, it requires human effort to filter out the false alignments. [40], [41] use automatic paraphrase models (e.g. fine-tuned BART [15]) to paraphrase clunky utterances and an automatic paraphrase filtering method to filter out low-quality paraphrases. The data synthesis method from [40], [41] requires the lowest paraphrase cost among all the aforementioned approaches. [23] generates synthetic data using SCFGs as well. [23] applies various approaches to down-sample a subset from the synthetic data. With their sampling method, training the parser with 200x less training data can perform comparably with training on the total population.

III. OUR PROPOSED APPROACH

Our proposed approach, as shown in Figure 1, improves on existing work for bootstrapping semantic parsers by not requiring manually annotated training data. Instead, large-scale training datasets can be generated in an automated manner through the use of existing automatic paraphrasing techniques. Specifically, we design a compact set of *synchronous grammar* rules to generate *seed* examples as pairs of canonical utterances and corresponding logical forms. We then apply a number of paraphrasing and filtering techniques to this initial set to create a much larger set of more diverse alternatives.

A. Semantic Parser

We adopt two common methods for training semantic parsers: i) the attentional SEQ2SEQ [20] framework which uses LSTMs [12] as the encoder and decoder, and ii) BERT-LSTM [39], a Seq2Seq framework with a copy mechanism [8] which uses RoBERTa [19] as the encoder and LSTM as the decoder. The input to the Seq2Seq model is the natural-language question and output is a sequence of linearized logical form’s tokens.

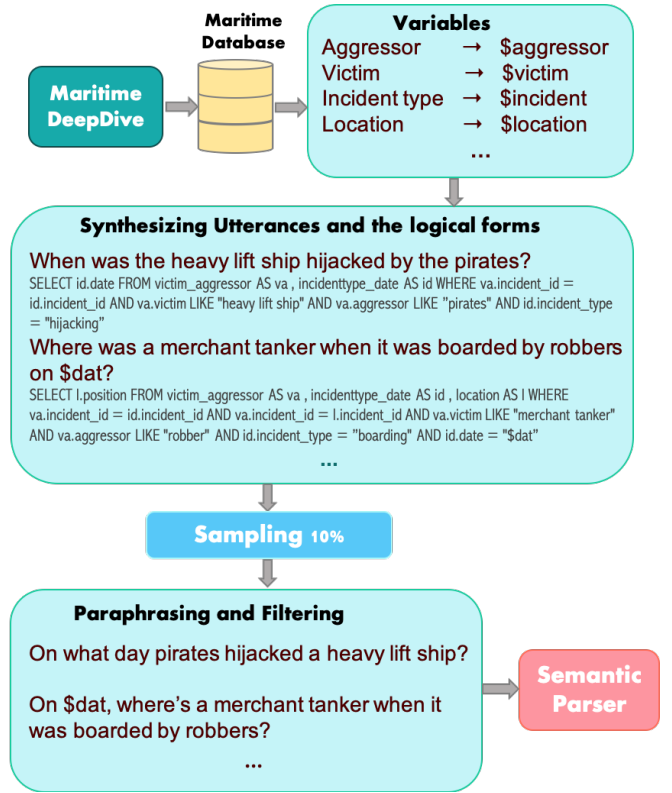


Fig. 1. An illustration of our end-to-end framework.

B. Synchronous grammar

Our constructed Maritime DeepDive Knowledge Graph [34], containing a set of entities such as *victim*, *aggressor* and triples such as (e_1, r, e_2) , where e_1 and e_2 are entities and r is a relationship (e.g., *victim_aggressor*). The database can be queried using SQL logical forms.

We propose to design compact grammars with only 31 SCFG rules that simultaneously generate both logical forms and canonical utterances such that the utterances are understandable by a human. First, we define variables specifying a canonical phrase such as *victim*, *aggressor*, *incident_type*, *date*, *position*, *location*. Second, we develop grammar rules for different SQL logical form structures (Figure 2). Finally, our framework uses the grammar rules and the list of variables and their possible values (G, L) to automatically generate canonical utterances paired with their SQL logical forms (u, lf) exhaustively (Figure 2). It yields a large set of canonical examples to train a semantic parser. We assign real values to the domain-specific variables including *victim* (victimized ships and individuals, e.g. oil tanker), *aggressor* (e.g. pirates), *incident_type* (e.g. robbery and hijacking). However, we replace all the general *location* (approximate place, e.g. country), *position* (place with longitude and latitude) and *date* with abstract variables $\$loc$, $\$pos$ and $\$dat$, which later can be used to capture the actual content using a NER model and parse into SQL queries. The location, position and date variables can have various unlimited content, and we do not restrict the

parser to some content.

C. Paraphrase Generation

The paraphrase generation model rewrites the canonical utterance to more diverse alternatives, which later are used to train the semantic parser. Existing work do not fully explore current paraphrasing approaches. Instead, they resort to applying language models (e.g., BART and GPT-3) to paraphrase utterances in datasets and do not tailor to specific domains.

We improve on the existing work and explore other augmentation methods for *paraphrase generation* on domain-specific data, making use of **back-translation**, **prompt-based language models**, **fine-tuned autoregressive language model** and **commercial paraphrasing model**.

1) *Paraphrasing using back-translation*: Back Translation is a data augmentation technique that is widely used in several studies [2], [25], [29]. Given a sentence, we aim to translate it to another language and translate it back, where the back-translated sentence will be slightly different from the original one. We thoroughly compared the performance of Google Translation API, a well-known translation tool for daily use. Google Translate API can translate 109 languages in all. It is supported by powerful neural machine translation methods enhanced with advanced techniques in a well-designed pipeline architecture.

To ensure the diversity of paraphrased data, we adopt the procedure as defined in [5]:

- Clustering languages into different family branches via Wikipedia info-boxes.
- Selecting the most used languages from each family and translating them using an appropriate translation system.
- Keeping the top three languages with the best translation performance.

2) *Paraphrasing using prompt-based language model generation*: Several studies [31]–[33] have investigated few-shot large pre-trained language models such as GPT-3 [3] for semantic parsing. Unlike the ‘traditional’ approach where a pre-trained language model can be leveraged by adapting its parameters to the task at hand through fine-tuning (*language model as pre-training task*), GPT-3 took a different approach where it can be treated as a ready-to-use multi-task solver (*language modelling as multi-task learning*). This could be achieved by transforming certain tasks we want it to solve into the form of language modelling. Specifically, by designing and constructing an appropriate input sequence of words (called a *prompt*), one is able to induce the model to produce the desired output sequence (i.e., a paraphrased utterance in this context) without changing its parameters through fine-tuning at all. In particular, in the low-data regime, empirical analysis shows that, either for manually picking hand-crafted prompts [21] or automatically building auto-generated prompts [8], [16] taking prompts for tuning models is surprisingly effective for the knowledge stimulation and model adaptation of pre-trained language model. However, none of these methods reports the use of a few-shot pre-trained language model to directly

generate few-shot and zero-shot paraphrased utterances. In this work, we supply GPT-3 with an instructive prompt to generate paraphrases in a zero-shot setting.

3) *Paraphrasing using fine-tuned language model generation*: Previous researches have proposed several pretrained autoregressive language models, such as GPT-2 [28] and ProphetNet [26]. Other studies such as [21] use these language models to solve various downstream tasks and achieve SOTA performance. In this work, following [41], we fine-tune an autoregressive language model BART [15] on the dataset from [14], which is a subset of the PARANMT corpus [38], to generate syntactically and lexically diversified paraphrases.

4) *Paraphrasing using a commercial system*: We also experiment with paraphrase generation using Quillbot.¹ Quillbot is a commercial tool that provides a scalable and robust paraphraser that can control synonyms and generation styles.

D. Paraphrase Filtering

Since the automatically generated paraphrases may have varying quality, we further filter out the paraphrases of low quality. In our work, we adopt the filtering method discussed in [40] in the spirit of self-training. The process consists of the following steps:

- Train the parser on the synthesized data generated by the SCFG rules.
- Evaluate the parser on the generated paraphrases and keep those for which the corresponding SQL logical forms are correctly generated by the parser.
- Add the paraphrase-SQL pairs into the training data and re-train the parser.
- Repeat steps 1–3 for several rounds or until no more paraphrases are kept.

This method is based on three assumptions [40]: i) the parser could generalize well to unseen paraphrases which share the same semantics with the original questions, ii) the synthetic dataset generated by the SCFGs are good enough to train an initial model, and iii) it is very unlikely for a poor parser to generate correct SQL queries by chance. To improve the model’s generalization ability, we use BERT-LSTM instead of vanilla Seq2Seq as our base parser for filtering.

In the following section, we provide an experimental analysis of the performance of our proposed approach.

IV. EXPERIMENTS

A. Maritime DeepDive

We utilize Maritime DeepDive [34], a probabilistic knowledge graph for the maritime domain automatically constructed from natural-language data collected from two main sources: (a) the Worldwide Threats To Shipping (WWTTS)² and (b) the Regional Cooperation Agreement on Combating Piracy and Armed Robbery against Ships in Asia (ReCAAP)³. We extract the relevant entities and concepts as well as their semantic

¹<https://quillbot.com>

²<https://msi.nga.mil/Piracy>

³<https://www.recaap.org/>

g_1 : what were \$aggressor armed with when \$incident \$victim \$date \$location ? g_2 : when did \$aggressor \$incident \$victim \$position ? g_3 : show me the list of \$incident \$date \$location . g_4 : what is \$superlative \$incident \$location \$date ?	u_1 : what were pirates armed with when attacking a product tanker on \$dat in \$loc ? u_2 : when did the armed gang hijack the underway tanker near position \$pos ? u_3 : show me the list of robberies on \$dat in \$loc . u_4 : what is the highest number of kidnapping in \$loc on \$dat ?
--	---

Fig. 2. Examples of utterances (right) generated from Synchronous grammar rules (left).

Canonical Utterance:	which weapon did pirates use to rob the offshore supply vessel on $\$dat$ in $\$loc$?
Techniques	Paraphrased Utterances
Back-translation (Spanish)	what weapon did the pirates used to steal the offshore supply ship in $\$dat$ at $\$loc$?
Back-translation (Telugu)	which weapon is used to rob the offshore supply vessel in $\$dat$ by pirates in $\$loc$?
Back-translation (Chinese)	which weapon is used in $\$loc$? it is used to rob the offshore supply boat?
GPT-3	what was the weapon used by pirates to rob the offshore supply vessel on $\$dat$ in $\$loc$?
BART	what gun did the pirates use to rob an offshore supply vessel in $\$loc$ on $\$dat$?
Quillbot	when pirates robbed an offshore supply vessel on $\$dat$ in $\$loc$, what weapon did they use?

TABLE I

EXAMPLES OF PARAPHRASING USING DIFFERENT TECHNIQUES.

relations, together with the uncertainty associated with the extracted knowledge. We consider the extracted maritime knowledge graph as our main database for building a QA system. Our training and test corpora include 1,235 and 217 piracy reports in the database, respectively.

B. Maritime Semantic Parsing Dataset

With 31 grammar rules, the list of variables and the Maritime database, we automatically synthesize 341,381 canonical utterances paired with SQL queries. As discussed above, there are semantic and grammatical mismatches between the synthetic canonical examples and real-world user-issued ones. Paraphrasing the synthesized dataset can potentially generate linguistically more diverse questions. Below, we test the paraphrasing strategies described in Section III to improve diversity of utterances while ensuring quality which yields the best end-to-end performance.

1) *Paraphrased Question Collection*: [23] shows that, with a proper sampling method, we can also achieve decent performance with much less synthetic training data. Following this insight, we only select a subset of synthetic canonical examples (10%) to reduce our paraphrase and training cost. We use a similar approach to Uniform Abstract Template Sampling (UAT) in [23] to get the diversified SQL logical forms structures for paraphrasing. There are 35,050 canonical examples in our sampled dataset, with 49% of samples containing abstract variables.

2) *SQL Query Writing and Validation*: We evaluate our semantic parser with real-world crowdsourced questions. Specifically, we collect 231 simple and complex questions about a set of given piracy reports. For all collected questions, we write their SQL queries manually and randomly divided them into a validation set with 76 examples and a test with 154 examples.

C. Paraphrase Generation

We employ and evaluate paraphrasing techniques with the following settings.

a) *Back-translation*: We choose three languages in Google Translation: Chinese (medium-resource), Spanish (high-resource) and Telugu (low-resource). The back-translated generation took approximately one day to complete the sampled data. In the initial paraphrased utterances, we found that Google Translation failed to distinguish the variable names in natural languages, which required the post-processing to correct these variable names. For example, variables like $\$pos$ will be translated to $\$POS$ or be removed in the paraphrased utterances. Another observation was that the performance of back-translated augmentation highly depends on the language, as shown in Table III.

b) *Prompt-based language model generation*: We use (GPT-3_Davinci) which is the most powerful model in the GPT-3 family. For GPT-3, generation quality can be boosted when more detailed and instructive prompts are given. Due to the limited data we have, we instruct GPT-3 to generate paraphrased questions in the zero-shot manner. Specifically, GPT-3 will directly generate relevant context by giving our manually-designed textual prompts. Different prompts will result in different outputs. The prompts and some generated examples can be found in Table II. The GPT-3 API from OpenAI generates paraphrase randomly based on the seed, and thus may have limited control on the generated output.

c) *Fine-tuned autoregressive language model generation*: We employ a pre-trained language model, (BART_Large) [15], to rewrite a canonical utterance u to a more diverse alternative. The BART_Large model⁴ is fine-tuned on a corpus

⁴We use the official implementation in fairseq, <https://github.com/pytorch/fairseq>.

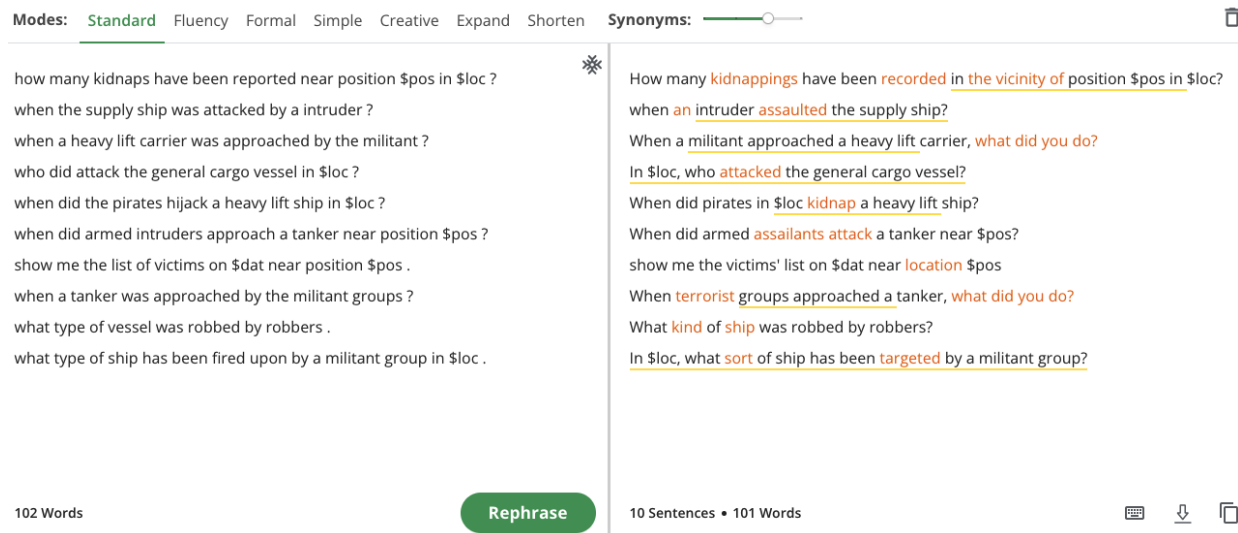


Fig. 3. We set the Quillbot paraphraser to the standard mode and the synonyms to have more changes.

Prompt: I'm a professional and creative paraphraser. I'm required to paraphrase the Original sentence in other words. The output sentence must not be the same as the given sentence.

Original: How many heavy lift vessels have been approached in \$loc ?
 Paraphrased: **How many heavy lift vessels have been contacted in \$loc?**

Prompt: I'm a professional and creative paraphraser.

Original: How many heavy lift vessels have been approached in \$loc ?
 Paraphrased: **How many heavy lift vessels have been approached in \$loc ?**

Prompt: I'm a professional and creative paraphraser. I'm required to paraphrase the Original sentence in other words.

Original: How many heavy lift vessels have been approached in \$loc ?
 Paraphrased: **How many heavy lift vessels have been contacted in your area?**

TABLE II

DIFFERENT PROMPTS RESULT IN DIFFERENT PARAPHRASE GENERATION. THE PARAPHRASED SENTENCE IN **BLUE** IS OUR EXPECTED OUTPUT, WHILE SENTENCES IN **PURPLE** ARE NOT IN THE IDEAL FORMAT.

of high-quality paraphrases sub-sampled from PARANMT [38] released by [14]. The model therefore learns to produce paraphrases with a variety of linguistic patterns, which is essential when paraphrasing from canonical utterances. The generated paraphrases are noisy or potentially vague. We reject paraphrases for which the parser cannot predict their SQL queries (Section IV-D).

d) Commercial paraphrasing model: As described in the previous section, we employ the commercial paraphrasing tool Quillbot to generate high-quality paraphrases. The setting of the paraphraser interface can be found in Figure 3.

D. Paraphrase Filtering

We develop a scheme where we input the paraphrased question into a parser, which in turn generates the SQL queries. If the SQL query matches exactly with the one generated from the original question of the paraphrase, we consider that this question has been correctly paraphrased.

Table III shows the filtering results of a number of paraphrasing techniques, including back-translation, fine-tuned BART, GPT-3 and Quillbot. In general, we see that the GPT-3 from OpenAI and Quillbot significantly outperform the other approaches. In contrast, the fine-tuned BART achieves only 20.22% retention. Moreover, we observe that GPT-3 has the highest semantic preserving ability by keeping 44.96% of the paraphrased sentences in the first round of filtering and Quillbot by keeps 51.91% and 54.25% in the second and third rounds of filtering suggesting prompt-based and commercial methods have good potential in paraphrasing. The Back-Translation method on Chinese performs the worst. We conjecture that Chinese is far from English in terms of linguistic structure; thus, Back-Translation suffers from translation errors. With more rounds of re-training, the generalization ability of the parser is improved, so more examples are kept. However, the improvement becomes marginal after the second round of re-training. Therefore, we only adopt three rounds since re-training is time-consuming.

E. Paraphrase Generation Method Analysis

Paraphrase diversity We measure the diversity of each paraphrase method by using the BLEU metric [24]. BLEU- N was originally proposed for automatic machine translation evaluation, which calculates the similarity score based on N -grams. The lower the score is, the more dissimilar the generated context is to the reference. The dissimilarity can imply the diversity of the paraphrased context when used in paraphrase evaluation. GPT-3 paraphrasing method gen-

#Rounds	BT (Spanish)	BT (Telugu)	BT (Chinese)	GPT-3	BART	Quillbot
Round 1	25.15	18.52	6.67	44.96	20.22	42.20
Round 2	29.63	37.85	13.78	51.81	22.34	51.91
Round 3	31.23	38.56	15.42	53.22	23.35	54.25
Total	30.21	40.91	16.37	53.37	19.68	54.92

TABLE III

% OF EXAMPLES KEPT AFTER EACH FILTERING ROUND.
 THE LOWEST % IN EACH ROUND IS HIGHLIGHTED IN BLUE.
 THE HIGHEST % IN EACH ROUND IS HIGHLIGHTED IN PURPLE.

Techniques	BT	GPT-3	*BART	Quillbot	All
Time (minutes)	8.67	33.3	0.43	60	102.4

TABLE IV

PRACTICAL PERFORMANCE (TIME COMPARISON PER 1000 SAMPLES).

*THE MEASUREMENT OF BART WAS RUN ON A SINGLE RTX 8000 NVIDIA GPU.

erates the least diversified sentences, while Back-Translation (Chinese) generates the most diversified ones. Quillbot and BART can generate paraphrases with a close diversity level. However, more Quillbot paraphrases are kept, indicating that Quillbot can perform well in terms of both diversity and the preservation of semantics. An interesting finding is that the results in Table V are relatively consistent with the filtering report in Table III. We assume it is easier for the parser to generalize on the paraphrases which have more words overlapping with the original questions.

Time-efficiency Table IV shows the required time for paraphrasing 1000 canonical utterances by each technique. The pre-trained BART model is the fastest way of paraphrasing while Quillbot is the slowest one.

Controllability. Among various paraphrasing techniques, Quillbot provides different settings and various paraphrased candidates (Figure 3). Current neural models still find it more challenging to get all possible replaced words or phrases while maintaining the semantic meaning.

Open-source. The second most promising paraphrasing model right now is provided by Quillbot. No other open-source model, except GPT-3, can achieve excellent quality compared to Quillbot. What techniques are practical for paraphrasing needs further investigation.

Compute resource. Unlike back-translation, where we can adopt translation systems to the paraphrasing tasks, direct paraphrasing needs specific paraphrasing data to train a good model. Designing a good paraphrasing dataset across different domains is challenging.

Readability. The back-translation system requires increasing human readability. Some translated sentences are not human-readable. A method should be proposed to automatically self-reject some invalid results. This will aid the back translation based paraphrasing method by providing more valid data.

F. Semantic Parser Evaluation

Baseline Models. We utilize two common methods of semantic parsers for the evaluation of different data settings:

i) the attentional SEQ2SEQ [20] framework which uses LSTMs [12] as the encoder and decoder, and ii) BERT-LSTM [39], a Seq2Seq framework with a copy mechanism [8] which uses RoBERTa [19] as the encoder and LSTM as the decoder. The inputs for these models are the natural-language questions and the outputs are their sequence of linearized SQL queries' tokens. Our evaluation metrics include Exact Matching as in [7], [18], Exact Matching (no-order), and Component Matching as in [42]

Component Matching: To conduct a detailed analysis of model performance, we measure the average exact match we are reporting F1 between the prediction and ground truth on different SQL components. For each of the following components: • SELECT • FROM • WHERE • GROUP BY • ORDER BY. We decompose each component in the prediction and the ground truth as bags of several sub-components, and check whether or not these two sets of components match exactly. In our evaluation, we treat each component as a set so that for example, WHERE va.aggessor = "pirates" AND va.victim = "container ship" and WHERE va.victim = "container ship" AND va.aggessor = "pirates" would be treated as the same query.

Exact Matching: We measure whether the predicted query as a whole is equivalent to the ground truth query.

Exact Matching (no-order): We first evaluate the SQL clauses and ignore the order in each component. The predicted query is correct only if all of the components are correct.

Table VI shows the performance of parsers on different training datasets. We have merged the filtered data using various paraphrasing techniques with the original training dataset and trained the semantic parsers with the new training dataset. We observe that Back-Translation paraphrase hurts the semantic parsing, indicating the generated sentences are too diversified and considered to be noisy signal during training. Among these Back-Translation methods, the model trained on filtered Back-Translation paraphrases with low-resource Telugu performs the worst. This may result from limited training data when training Google Translation on Telugu. We also find that paraphrasing only 10% of the original training dataset using GPT-3 and Quillbot can constantly boost the semantic parsing. We further demonstrate in Table VII that with more Quillbot paraphrase data, the performance of semantic parser keeps increasing. With the findings in Tables III and V, we conclude paraphrase filtering and corresponding diversity are

Metrics	BT (Spanish)	BT (Telugu)	BT (Chinese)	GPT-3	BART	Quillbot
BLEU-1	59.05	48.97	36.95	74.81	53.39	59.74
BLEU-2	43.65	36.26	23.89	68.42	40.32	51.22
BLEU-3	33.22	26.60	15.98	63.48	31.12	44.63
BLEU-4	25.63	19.33	10.30	59.46	24.17	39.55

TABLE V

DIVERSITY EVALUATION ON PARAPHRASE METHODS.

THE LOWEST SCORE IN EACH BLEU METRIC IS HIGHLIGHTED IN BLUE.

THE HIGHEST SCORE IN EACH BLEU METRIC IS HIGHLIGHTED IN PURPLE.

Training Data	SEQ2SEQ			RoBERTa-base		
	Exact match (acc)	Exact match no order (acc)	Compo_match (F1)	Exact match (acc)	Exact match no order (acc)	Compo_match (F1)
Original dataset	36.77	48.39	70.71	43.07	57.79	77.26
BT (Spanish)	37.42	50.97	75.09	42.86	57.14	77.11
BT (Chinese)	36.77	54.84	74.61	45.44	59.44	78.17
BT (Telugu)	34.84	49.68	73.19	42.86	57.14	76.72
BART	36.77	54.84	75.60	45.02	59.93	78.93
GPT-3	43.87	58.86	76.79	47.05	61.64	79.04
Quillbot	41.29	55.48	75.53	50.65	62.99	80.58
Full Filtered Data	40.00	56.13	76.19	46.75	60.17	78.55

TABLE VI

EVALUATION OF THE SEMANTIC PARSER BASELINES BASED ON ADDITIONAL PARAPHRASED DATA USING DIFFERENT TECHNIQUES.

Training Data	Filtering	Compo_match (F1)
Original training dataset	-	77.26
With Quillbot (10%)	54.25	80.58
With Quillbot (20%)	64.73	81.73

TABLE VII

THE EFFECT OF DATA PARAPHRASING PORTION ON FILTERING AND PARSER (ROBERTA-BASE) LOGICAL FORM MATCHING ACCURACY

not directly related to the semantic parsing performance. A more comprehensive and efficient approach needs to be found to evaluate paraphrase quality in terms of semantic parsing.

V. OBSERVATIONS AND DISCUSSIONS

By conducting several experiments and using the obtained results on semantic parsers, we now present our observations and insights on the task.

- With the finding in Table VI we observe that paraphrasing is a powerful technique to improve the performance of the final QA system without additional human efforts.
- Table VII shows that with only 10% additional paraphrased questions (using Quillbot) we are able to increase the average F1 by 3.32% and with further 10% additional paraphrased questions, F1 increases by 4.47%. We can conclude that a small portion of high quality paraphrased data would be sufficient for training semantic parsers.
- Among languages we choose for Back-Translation paraphrasing in TABLE V, Chinese shows the lowest BLEU score and the highest diverse paraphrased questions. Subsequently, the semantic parser (RoBERTa-base) demonstrates a better performance with Chinese compare to the other two languages (i.e. Spanish and Telugu). Further

analyses is required to investigate the impact of amount of filtering data and diversity on parsers' performance.

- In the last row of TABLE VI, we simply combined the all paraphrased data from all paraphrasing techniques. In the future, we will investigate whether combining the data from only more effective paraphrasing techniques, such as Quillbot and GPT-3, can induce more promising results.

VI. CONCLUSION

This work is aimed at enhancing human-machine communications by automatically translating users' questions expressed in natural language into SQL queries. By developing a compact set of synchronous grammar rules enhanced with a set of multiple automatic paraphrasing and filtering techniques in order to generate large scale training data, we are able to train a semantic parser without the need to use any (manually) annotated training dataset. The experimental results demonstrated promising results could be achieved at a very low cost, which is important in many real-world situations where data is scarce and annotated data is not available.

REFERENCES

- [1] Laura Banarescu, Claire Bonial, Shu Cai, Madalina Georgescu, Kira Griffitt, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Martha Palmer, and Nathan Schneider. Abstract meaning representation for sembanking. In *Proceedings of the 7th linguistic annotation workshop and interoperability with discourse*, pages 178–186, 2013.
- [2] Djamila Romaiissa Beddiar, Md Saroar Jahan, and Mourad Oussalah. Data expansion using back translation and paraphrasing for hate speech detection. *Online Social Networks and Media*, 24:100153, 2021.
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

- [4] Bo Chen, Le Sun, and Xianpei Han. Sequence-to-action: End-to-end semantic graph generation for semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 766–777, 2018.
- [5] Jean-Philippe Corbeil and Hadi Abdi Ghadivel. Bet: A backtranslation approach for easy data augmentation in transformer-based paraphrase identification context. *arXiv preprint arXiv:2009.12452*, 2020.
- [6] Li Dong and Mirella Lapata. Language to logical form with neural attention. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33–43, 2016.
- [7] Li Dong and Mirella Lapata. Coarse-to-fine decoding for neural semantic parsing. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 731–742, 2018.
- [8] Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, 2016.
- [9] Jiaqi Guo, Zecheng Zhan, Yan Gao, Yan Xiao, Jian-Guang Lou, Ting Liu, and Dongmei Zhang. Towards complex text-to-sql in cross-domain database with intermediate representation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4524–4535, 2019.
- [10] Sonal Gupta, Rushin Shah, Mrinal Mohit, Anuj Kumar, and Mike Lewis. Semantic parsing for task oriented dialog using hierarchical representations. *arXiv preprint arXiv:1810.07942*, 2018.
- [11] Jonathan Herzig and Jonathan Berant. Don’t paraphrase, detect! rapid and effective data collection for semantic parsing. *arXiv preprint arXiv:1908.09940*, 2019.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [13] Aishwarya Kamath and Rajarshi Das. A survey on semantic parsing. In *Automated Knowledge Base Construction (AKBC)*, 2018.
- [14] Kalpesh Krishna, John Wieting, and Mohit Iyyer. Reformulating unsupervised style transfer as paraphrase generation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, 2020.
- [15] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020.
- [16] Zhuang Li, Lizhen Qu, and Gholamreza Haffari. Context dependent semantic parsing: A survey. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2509–2521, 2020.
- [17] Zhuang Li, Lizhen Qu, and Gholamreza Haffari. Total recall: a customized continual learning method for neural semantic parsers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3816–3831, 2021.
- [18] Zhuang Li, Lizhen Qu, Shuo Huang, and Gholamreza Haffari. Few-shot semantic parsing for new predicates. *arXiv preprint arXiv:2101.10708*, 2021.
- [19] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [20] Minh-Thang Luong, Hieu Pham, and Christopher D Manning. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- [21] Farhad Moghimifar, Lizhen Qu, Yue Zhuo, Mahsa Baktashmotlagh, and Gholamreza Haffari. Cosmo: Conditional seq2seq-based mixture model for zero-shot commonsense question answering. *arXiv preprint arXiv:2011.00777*, 2020.
- [22] Van Nguyen and Liam Mellor. Fuzzy mlms and qstags for activity recognition and modelling with rush. In *2020 IEEE 23rd International Conference on Information Fusion (FUSION)*, pages 1–8. IEEE, 2020.
- [23] Inbar Oren, Jonathan Herzig, and Jonathan Berant. Finding needles in a haystack: Sampling structurally-diverse training sets from synthetic data for compositional generalization. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10793–10809, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [24] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July 2002. Association for Computational Linguistics.
- [25] Shrimai Prabhumoye, Yulia Tsvetkov, Ruslan Salakhutdinov, and Alan W Black. Style transfer through back-translation. *arXiv preprint arXiv:1804.09000*, 2018.
- [26] Weizhen Qi, Yu Yan, Yeyun Gong, Dayiheng Liu, Nan Duan, Jiusheng Chen, Ruofei Zhang, and Ming Zhou. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training. *arXiv preprint arXiv:2001.04063*, 2020.
- [27] Maxim Rabinovich, Mitchell Stern, and Dan Klein. Abstract syntax networks for code generation and semantic parsing. *arXiv preprint arXiv:1704.07535*, 2017.
- [28] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [29] Djamila Romaiha Beddiar, Md Saroar Jahan, and Mourad Oussalah. Data expansion using back translation and paraphrasing for hate speech detection. *arXiv e-prints*, pages arXiv–2106, 2021.
- [30] Adam Saulwick. Lexpresso: a controlled natural language. In *International Workshop on Controlled Natural Language*, pages 123–134. Springer, 2014.
- [31] Nathan Schucher, Siva Reddy, and Harm de Vries. The power of prompt tuning for low-resource semantic parsing. *arXiv preprint arXiv:2110.08525*, 2021.
- [32] Richard Shin, Christopher H Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. Constrained language models yield few-shot semantic parsers. *arXiv preprint arXiv:2104.08768*, 2021.
- [33] Richard Shin and Benjamin Van Durme. Few-shot semantic parsing with language models trained on code. *arXiv preprint arXiv:2112.08696*, 2021.
- [34] Fatemeh Shiri, Van Nguyen, Shuang Yu, Teresa Wang, Shirui Pan, Xiaojun Chang, Yuan-Fang Li, and Reza Haffari. Toward the automated construction of probabilistic knowledge graphs for the maritime domain. In *2021 IEEE 24th International Conference on Information Fusion (FUSION)*, pages 1–8. IEEE, 2021.
- [35] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27, 2014.
- [36] Bailin Wang, Richard Shin, Xiaodong Liu, Oleksandr Polozov, and Matthew Richardson. 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, 2020.
- [37] Yushi Wang, Jonathan Berant, and Percy Liang. Building a semantic parser overnight. 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 1332–1342, 2015.
- [38] John Wieting and Kevin Gimpel. Parantmt-50m: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, 2018.
- [39] Silei Xu, Giovanni Campagna, Jian Li, and Monica S Lam. Schema2qa: High-quality and low-cost q&a agents for the structured web. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1685–1694, 2020.
- [40] Silei Xu, Sina J Semnani, Giovanni Campagna, and Monica S Lam. Autoqa: From databases to qa semantic parsers with only synthetic training data. *arXiv preprint arXiv:2010.04806*, 2020.
- [41] Pengcheng Yin, John Wieting, Avirup Sil, and Graham Neubig. On the ingredients of an effective zero-shot semantic parser. *arXiv preprint arXiv:2110.08381*, 2021.
- [42] Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *EMNLP*, 2018.
- [43] Q. Zhu, X. Ma, and X. Li. Statistical learning for semantic parsing: A survey. *Big Data Mining and Analytics*, 2(4):217–239, Dec 2019.