

Semantic Structure based Query Graph Prediction for Question Answering over Knowledge Graph

Mingchen Li

Department of Computer Science
Georgia State University, Atlanta, USA
mli33@student.gsu.edu

Shihao Ji

Department of Computer Science
Georgia State University, Atlanta, USA
sji@gsu.edu

Abstract


Building query graphs from natural language questions is an important step in complex question answering over knowledge graph (Complex KGQA). In general, a question can be correctly answered if its query graph is built correctly and the right answer is then retrieved by issuing the query graph against the KG. Therefore, this paper focuses on query graph generation from natural language questions. Existing approaches for query graph generation ignore the semantic structure of a question, resulting in a large number of noisy query graph candidates that undermine prediction accuracies. In this paper, we define six semantic structures from common questions in KGQA and develop a novel Structure-BERT to predict the semantic structure of a question. By doing so, we can first filter out noisy candidate query graphs, and then rank the remaining candidates with a BERT-based ranking model. Extensive experiments on two popular benchmarks MetaQA and WebQuestionsSP (WSP) demonstrate the effectiveness of our method as compared to state-of-the-arts. The source code can be found at <https://github.com/ToneLi/SSKGQA>.

1 Introduction

Knowledge graph (KG) is a graph structured database (Miller, 1995), in which nodes represent entities (e.g., *Hedgehog in the Fog*, *Sergei Kozlov*), and edges reflect the relations between entities (e.g., *Hedgehog in the Fog* - *written_by* - *Sergei Kozlov*). Users can get crisp answers by querying KGs with natural language questions, and this process is called Question Answering over Knowledge Graph (KGQA). Recently, consumer market witnesses a widespread application of this technique in a variety of virtual assistants, such as Apple Siri, Google Home, Amazon Alexa, and Microsoft Cortana, etc.

Early works (Bordes et al., 2015; Golub and He, 2016) on KGQA mainly focus on simple ques-

Question 1: Which person wrote the films directed by Yuriy Norshteyn?



Question 2: What is Cher's son's name?




Figure 1: (Q1) Example question involving multi-hop reasoning, and (Q2) Example question with constraints

tions, such as *where is the hometown of Obama?* This question only involves in one relation path (e.g., *hometown* or *birth-place*) in KG, and is relatively easy to solve. However, many questions in daily QA sessions are often more complex, manifested by multi-hop reasoning or questions with multiple constraints. Therefore, recently there is a flurry of interests on complex KGQA (Shi et al., 2021; Yadati et al., 2021).

There are two types of complexity when dealing with complex KGQA, i.e., multi-hop questions and questions with multiple constraints (See Figure 1 for example). Question 1 in Figure 1 is a typical multi-hop question, to which the answer is related to Yuriy Norshteyn with two-hop relations: *directed.by.reverse* and *written.by*. In response to this challenge, Xu et al. (2019) enhances the traditional Key-Value Memory Neural Networks (KV-MemNNs) (Miller et al., 2016) for multi-hop question answering. They design a query updating strategy to decompose the question and predict the relevant relation paths at each hop. TransferNet (Shi et al., 2021) is an effective and transparent model for multi-hop questions; it attends to different parts of the question at each hop and computes activated scores for relation path prediction. Despite the promising results, it's still challenging for these models to

predict relation paths accurately at each hop, and thus suffer from error propagation over multi-hop reasoning. Similarly, Question 2 in Figure 1 is an example of question with constraints. Apparently, there is a single relation path between the topic entity *Cher* and the answer *Chaz Bono*, but the constraint of `person.gender=Male` must be satisfied. To handle this type of complex questions, many works built on the idea of query ranking are proposed (Yih et al., 2015; Lan and Jiang, 2020; Chen et al., 2020), which rank candidate queries by the similarity scores between question and candidate queries. Specifically, these ranking methods use query graphs to represent queries, and explore various strategies to generate candidate query graphs for ranking. Typical strategies assume the answers are within n hops of topic entity, and enumerate all the relation paths within n hops to generate candidate query graphs. Although this candidate generation strategy can yield all valid query graphs from a topic entity, they have two main limitations: (1) The generated candidate query graphs are very noisy. As shown in Figure 2(a), a candidate query graph with an incorrect structure is presented; this candidate query graph is generated by the traditional enumeration strategy but lacks of the constraint on `person.gender`, which can incur considerable error in query graph ranking (See Table 4). For the example in Figure 2(a), both `parent` and `birth.place` are relevant to the question; even though this candidate query graph has an incorrect semantic structure (to be defined in Sec. 3.2), it is still challenging for ranking models to demote it below the correct query graph – the one in Figure 2(b). (2) When building a ranking model to rank query graphs, recent works (Lan and Jiang, 2020) treat the candidate query graph and question as a sequence of words and leverage BERT (Devlin et al., 2018) to extract feature representation from its pooled output. However, this pooled output is usually not a good representation of the semantics of the input sequence (Khodeir, 2021). Therefore, improved ranking models are to be developed.

To mitigate the aforementioned issues, this paper proposes *SSKGQA*, a Semantic Structure based framework for complex KGQA. We represent both the multi-hop questions and questions with constraints in a way similar to query graphs¹ and rank the generated candidate query graphs by the simi-

¹For example, one candidate query graph for Question 1 in Figure 1 can be written as $Yuriy.Norshhteyn \rightarrow directed.by.reverse \rightarrow y \rightarrow written.by \rightarrow x$.

Question: What is the birth place of Queen Isabella 's mother?

Answer: Lisboa




Figure 2: Example two-hop candidate query graphs for a question with constraints. (a) A candidate query graph with an incorrect semantic structure, (b) A candidate query graph with a correct semantic structure.

ilarity scores between question and candidate query graphs. Inspired by Chen et al. (2020), if the structure of a question is known in advance, the noise in candidate query graphs can be reduced significantly by filtering. Thus, *SSKGQA* first predicts the semantic structure of a natural language question, which is then used to filter out noisy candidate query graphs (which have incorrect structures).

Specifically, we define six semantic structures based on the question topology that is introduced by Srivastava et al. (2021). With the defined semantic structures, our *SSKGQA* processes a natural language question in two stages. In the first stage, we develop a novel Structure-BERT to predict the semantic structure of a natural language question, which is then used to filter out noisy candidate query graphs and produce a set of query graph candidates that match the predicted structure. In the second stage, we rank the remaining candidate query graphs of a question by a BERT-based ranking model and identify the top-1 candidate query graph, which is then issued to retrieve the final answer from a KG. Our experiments demonstrate that this semantic structure based query graph prediction strategy is very effective and enables *SSKGQA* to outperform state-of-the-art methods.

Our main contributions are summarized as follows. (1) We propose *SSKGQA*, a semantic structure based method to predict query graphs from natural language questions. *SSKGQA* can handle both multi-hop questions and questions with constraints and is a unified framework for complex KGQA. (2) We develop a novel Structure-BERT to predict the semantic structure of each question, and a BERT-based ranking model with a triplet loss to identify the top-1 query graph candidate. (3) Compared to state-of-the-arts methods, our *SSKGQA* demonstrates superior performance on two popular complex KGQA benchmarks.

2 Related Work

2.1 Multi-hop Question Answering

Current works on multi-hop question answering mainly focus on how to retrieve answers by calculating the relation paths step by step. In general, a right answer can be retrieved if the relation paths are identified correctly at each hop. Xu et al. (2019) enhances the traditional Key-Value Memory Neural Networks (KV-MemNNs) (Miller et al., 2016) and designs a query updating strategy to decompose the question and predict the relevant relation paths at each hop. TransferNet (Shi et al., 2021) calculates the relation path scores based on an updated question at each hop, but they leverage the attention mechanism to update question representations over multiple hops. More recently, Cai et al. (2021) introduces the dual process theory to predict the relation paths at each hop. Although these methods achieve promising results, they suffer from error propagation when predicting the relation paths step by step. To mitigate this issue, SSKGQA identifies the top-1 query graph by directly calculating the similarity scores between question and candidate query graphs (or similarly relation paths).

On the other hand, Sun et al. (2019a, 2018) incorporate external corpus to enhance the performance of KGQA. They focus on how to get the answers by constructing a subgraph for each question. A challenge of this method is that it is difficult to construct a subgraph around topic entity because we need to identify relevant entities from external corpus, and this process is error-prone. Saxena et al. (2020) predicts the answers by utilizing the KG embedding model. However, complex questions with long relation paths can reduce the learnability of KG embedding significantly. Our SSKGQA does not need external corpus to improve prediction accuracies and can solve complex multi-hop questions by a semantic structure based ranking.

2.2 Complex Questions with Constraints

For questions with constraints, a sequence of works focus on how to reach the answers by generating query graphs. Yih et al. (2015) enumerates all possible entities and relation paths that are connected to a topic entity to generate candidate query graphs, and uses a CNN-based ranking model to identify the query graph. Following a similar candidate query graph generation of (Yih et al., 2015), Maheshwari et al. (2019) propose a novel query graph ranking method based on self-attention. Qin et al.

(2021) introduces a query graph generation method by using their proposed relation subgraphs. However, these methods largely ignore the noise when generating the candidate query graphs, which undermines the predictive performance during query graph ranking. To mitigate this issue, SSKGQA first predicts the semantic structure of a question, which is then used to reduce the noise in candidate query graphs.

2.3 Query Graph Ranking

Current research on KGQA mainly focuses on how to generate the candidate query graphs, and there are only a few works exploring how to rank the candidate query graphs. Lan and Jiang (2020) concatenates question and candidate query graph into a single sequence, and leverages BERT (Devlin et al., 2018) to process the whole sequence for ranking. However, Reimers and Gurevych (2019) show that this strategy is inefficient as it can incur a massive computation due to the combinatorial nature of concatenation of question and candidate query graphs, leading to duplicated calculation. Chen et al. (2020) explore GRUs to encode the question and query graph information, and utilize a hinge loss to learn a ranking model. However, GRUs can only learn a limited interaction among words in a sentence, while the global interactions among words has proven to be critical for text representation in various NLP applications (Khan et al., 2020). To solve the aforementioned issues, SSKGQA exploits separated BERT models to process questions and query graphs, respectively, and reuses the extracted features to avoid duplicated calculation and leverages a triplet loss to train a BERT-based ranking model.

3 Preliminaries

3.1 Query Graph

Query graph is a graph representation of a natural language question (Yih et al., 2015). See Figure 2 for example. A query graph usually contains four types of components: (1) a grounded entity, which is an entity in KG and is often the topic entity of a question, e.g., Queen Isabella in Figure 2. (2) an existential variable, which is an ungrounded entity, e.g., Isabel de Portugal in Figure 2. (3) a lambda variable, which is the answer to question but usually an ungrounded entity, e.g., Lisboa in Figure 2. (4) some constraints on a set of entities, e.g., gender in Figure 2. A ques-




Figure 3: Six semantic structures defined in the paper. There are three semantic structures for questions in MetaQA: ($SS1$, $SS2$, $SS3$), and five semantic structures for questions in WSP: ($SS1$, $SS2$, $SS4$, $SS5$, $SS6$).

tion can be correctly answered if its query graph is built correctly and the right answer can be retrieved by issuing the query graph (represented by a SPARQL (Pérez et al., 2009) command) to a KG.

3.2 Semantic Structures

As observed by (Chen et al., 2020), if the structure of a question is known in advance, the noise in candidate query graphs can be reduced significantly by filtering. Thus, in this paper we define six semantic structures based on the question topology that is introduced by Srivastava et al. (2021). These six semantic structures are listed in Figure 3 and example questions for each semantic structure can be found in Figure 6 in Appendix. As we can see, a semantic structure is a graph that is an *abstract* of the query graphs of the same pattern. Typically, a semantic structure consists of four components $\{E, r, v, C\}$, where E denotes an entity, r refers to all types of relations, v is an existential variable, and C denotes a constraint.

To identify the semantic structure of a question, we can train a classifier for prediction. But first we need to annotate each training question with its semantic structure. Fortunately, this annotation can be achieved readily for questions in MetaQA and WebQuestionsSP (WSP) since these questions are either partitioned by number of hops or accompanied by the SPARQL commands. Details on question annotation are provided in Sec. 7.2 in Appendix. By annotating the questions in MetaQA and WSP, we found that these six semantic structures can cover 100% of questions in MetaQA, and 77.02% of questions in the test set of WSP as shown in Table 1. It is challenging to design additional semantic structures to cover 100% of questions in WSP because there are some unusual operators in WSP, such as Or and \leq , which are difficult to map to a common semantic structure. Even though there is only a 77.02% coverage on the WSP test questions, our experiments show that SSKGQA already

| Dataset | MetaQA | | | WSP |
|---------|--------|-------|-------|-------|
| | Hop-1 | Hop-2 | Hop-3 | |
| Train | 100 | 100 | 100 | 91.37 |
| Test | 100 | 100 | 100 | 77.02 |
| Dev | 100 | 100 | 100 | n/a |

Table 1: Semantic structure coverage (%) for questions in the training, test and development sets of MetaQA and WSP.

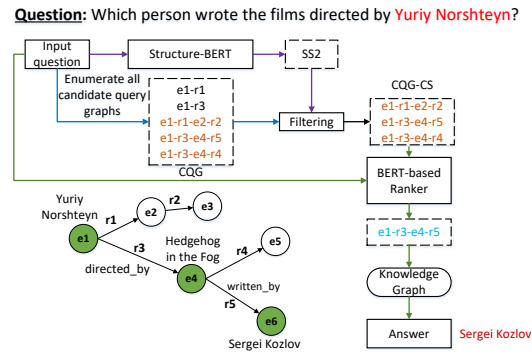


Figure 4: Overview of SSKGQA. A subgraph related to Yuriy Norshteyn from a KG is provided for illustration.

outperforms state-of-the-art methods on WSP. As a future work, we plan to explore new techniques to cover the remaining 22.98% of questions in WSP.

4 The Proposed Method

We first provide an overview of SSKGQA, and then discuss its main components: (1) Structure-BERT and (2) query graph ranking in details.

4.1 Overview

The overview of our proposed SSKGQA is depicted in Figure 4. Given a question q , following previous works (Saxena et al., 2020; Chen et al., 2020; Cai et al., 2021) we assume the topic entity of q has been obtained by preprocessing. Then the answer to q is generated by the following steps. First, the semantic structure of q is predicted by a novel Structure-BERT classifier. For the example in Figure 4, q is a 2-hop question and the classifier predicts its semantic structure as $SS2$. Second, we retrieve all the candidate query graphs (CQG) of q by enumeration², and use the predicted semantic structure $SS2$ as the constraint to filter out

²For clarify, only 1-hop and 2-hop candidate query graphs are considered in this example.

noisy candidate query graphs and keep the candidate query graphs with correct structure (CQG-CS). Afterwards, a BERT-based ranking model is used to score each candidate query graph in CQG-CS, and the top-1 highest scored candidate is selected as the query graph g for question q . Finally, the selected query graph is issued to KG to retrieve the answer Sergei Kozlov.

4.2 Structure-BERT

Given a question q , we first need to predict its semantic structure, which is a multi-class classification problem that classifies question q to one of the six semantic structures defined in Figure 3.




Figure 5: Structure-BERT: given a question and its topic entity, the model predicts the semantic structure of the question.

Figure 5 depicts the architecture of Structure-BERT. The input to Structure-BERT is question q and its topic entity te . The output of Structure-BERT is a probability distribution over six semantic structures, i.e., $p(y|q, \theta)$, where θ denotes the model parameters of Structure-BERT. Structure-BERT contains three sub-modules. **Question Encoder** encodes question q by a BERT language model (Devlin et al., 2018), and the final hidden state corresponding to token [CLS] is used as the question embedding e^q . **Entity Encoder** leverages a pre-trained knowledge embedding model, such as TransE (Bordes et al., 2013) or ComplEx (Trouillon et al., 2016), to extract the entity embedding e^h . Next, the extracted question embedding e^q and entity embedding e^h are fed to a **RotatE** module for information fusion. First, we utilize a pre-trained RotatE (Sun et al., 2019b) model to calculate a “tail” embedding e^t using e^h and e^q , and then fuse the topic entity, question and “tail” embeddings by combining e^h , e^q and e^t to a latent representation s by

$$\begin{aligned} \bar{e}^h &= e_{lh}^h + ie_{hh}^h \\ \bar{e}^q &= e_{lh}^q + ie_{hh}^q \\ \bar{e}^t &= \bar{e}^h \times \bar{e}^q = e_{lh}^h e_{lh}^q - e_{hh}^h e_{hh}^q + i(e_{hh}^h e_{lh}^q + e_{lh}^h e_{hh}^q) \\ s &= e^h + e^q + e^t, \end{aligned} \quad (1)$$

where e_{lh}^* (e_{hh}^*) denotes the lower (higher) half of vector e^* . As such, \bar{e}^* is a complex vector whose real (imaginary) part is from e_{lh}^* (e_{hh}^*). Therefore, we can convert between e^* and \bar{e}^* readily.

Finally, the latent representation s is fed to a fully connected layer, followed by a softmax for classification. The whole network is fully differentiable and can be optimized by minimizing the traditional cross-entropy loss.

To train Structure-BERT, we need to annotate questions with their semantic structures to develop a training, validate and test set. As discussed in Sec. 3.2, this annotation can be conducted readily for questions in MetaQA and WSP. The details are provided in Sec. 7.2 in Appendix.

4.3 Query Graph Ranking

Another important component of SSKGQA is a BERT-based ranking model for query graph ranking that can be trained by a triplet loss (Schroff et al., 2015). Specifically, the ranking model has three inputs: (1) question $q = \{ [\text{CLS}], w_1, w_2, \dots, w_m, [\text{SEP}] \}$, where w_i is the i -th word of q ; (2) positive query graph³ $g^p = \{ [\text{CLS}], u_1, u_2, \dots, u_n, [\text{SEP}] \}$, where u_i is the i -th unit of a query graph that is split by space or special symbol such as “.”. For example, given query graph (Natalie Portman, film.actor.film, v), u_1 =Natalie, u_2 =Portman, ... and u_6 = v ; and (3) negative query graph g^n that is a candidate query graph in CQG-CS except the positive candidate of the question.

We utilize a BERT model $f(\cdot)$ to extract the semantic representations $f(q)$, $f(g^p)$, $f(g^n)$ for q , g^p , g^n , respectively. This BERT model is built on a pre-trained BERT from Hugging Face⁴; we add one extra multi-head attention layer on top of the hidden state of the pre-trained BERT (See the ablation study in Sec. 5). This BERT-based ranking model $f(\cdot)$ is then optimized by minimizing the

³The positive query graph of a question can be found from the relation paths provided in MetaQA or the SPARQL command provided in WSP.

⁴<https://huggingface.co/bert-base-uncased>

triplet loss (Schroff et al., 2015)

$$\max(\|f(q)-f(g^p)\|-\|f(q)-f(g^n)\|+\alpha,0), \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean distance and α is a margin parameter, which we set to 1 as default.

During training, the triplet loss reduces the distance between $f(q)$ and $f(g^p)$, while enlarging the distance between $f(q)$ and $f(g^n)$. At inference time, we calculate the similarity scores between question and its candidate query graphs from CQG-CS, and choose the top-1 highest scored candidate as query graph g to retrieve final answer from KG.

5 Experiments

We evaluate the performance of SSKGQA on two popular KGQA benchmarks: MetaQA and WebQuestionsSP (WSP), and compare it with seven state-of-the-arts methods. Ablation study is also conducted to understand the effectiveness of different components of SSKGQA.

Our PyTorch source code is provided at <https://github.com/ToneLi/SSKGQA>. All our experiments are performed on Nvidia RTX GPUs.

5.1 Datasets

- **MetaQA** (Zhang et al., 2018) is a large scale KGQA dataset with more than 400k questions. It contains questions with 1, 2 or 3 hops. In our experiments, we use the vanilla version of the QA dataset. MetaQA also provides a KG from the movie domain with 43,233 entities, 9 relations and 134,741 triples.
- **WebQuestionsSP (WSP)** (Yih et al., 2016) is a small scale KGQA dataset with 5,119 questions which are answerable through Freebase KG. Since Freebase has more than 338,580,000 triples, for ease of experimentation we use a light version provided by Saxena et al. (2020). This smaller KG has 1.8 million entities and 5.7 million triples.

The statistics of training, development and test datasets of MetaQA and WSP is provided in Table 2. Compared to MetaQA, WSP is relatively small QA dataset even though its KG is much larger than that of MetaQA.

5.2 Hyperparameter Settings

Structure-BERT We set the dropout rate to 0.1, batch size to 32, and use AdamW optimizer (Loshchilov and Hutter, 2017) with the learning rate of $5e-8$. We also apply gradient clipping

| Dataset | Train | Dev | Test |
|--------------|---------|--------|--------|
| MetaQA- hop1 | 96,106 | 9,992 | 9,947 |
| MetaQA- hop2 | 118,980 | 14,872 | 14,872 |
| MetaQA- hop3 | 114,196 | 14,274 | 14,274 |
| WSP | 3,304 | – | 1,815 |

Table 2: Statistics of the MetaQA and WSP datasets

to constrain the maximum value of L_2 -norm of the gradients to be 1. To extract the latent representations of topic entities, pre-trained ComplEx (Trouillon et al., 2016) and TransE (Bordes et al., 2013) are adopted for MetaQA and WSP, respectively.

BERT-based Ranking Model We add one extra multi-head attention layer on top of the hidden state of the pre-trained BERT. This extra multi-head attention layer contains three attention heads and a 3072-dim fully connected layer. The dropout rate is set to 0.5. We use AdamW Optimizer (Loshchilov and Hutter, 2017) with the learning rate of $2e-5$. We also use gradient clipping to constrain the max L_2 -norm of the gradients to be 1.

5.3 Baselines

We compare our SSKGQA against seven state-of-the-art complex KGQA models: 1) GraftNet (Sun et al., 2018), which answers the questions based on the subgraphs it creates. 2) PullNet (Sun et al., 2019a), which proposes a “pull” operation to retrieve the relevant information from KG and external corpus. 3) Key-Value Memory Network (KV-MemNN) (Miller et al., 2016), which uses key-value pairs as the memory units to answer questions. 4) EmbedKGQA (Saxena et al., 2020), which proposes a knowledge embedding method for Complex KGQA. 5) TransferNet (Shi et al., 2021), which utilizes an interpretable model for complex KGQA. 6) DCRN (Cai et al., 2021), which proposes a Bayesian network to retrieve the final answers. For MetaQA, we also include 7) VRN (Zhang et al., 2018) as the baseline, which proposes an embedding reasoning graph and utilizes variational inference to improve the performance of Complex KGQA.

5.4 Comparison with State-of-the-Arts

Table 3 reports the performances of SSKGQA and seven state-of-the-art methods on MetaQA and WSP. As can be seen, the performances of KV-

| Model | Hop-1 | Hop-2 | Hop-3 | WSP |
|-------------|-------------|------------|------------|-------------|
| KV-MemNN | 96.2 | 82.7 | 48.9 | 46.7 |
| VRN | 97.5 | 89.2 | 62.5 | - |
| GraftNet | 97.0 | 94.8 | 77.7 | 66.4 |
| PullNet | 97.0 | 99.9 | 91.4 | 68.1 |
| EmbedKGQA | 97.5 | 98.8 | 94.8 | 66.1 |
| DCRN | 97.5 | 99.9 | 99.3 | 67.8 |
| TransferNet | 97.5 | 100 | 100 | 71.4 |
| SSKGQA | 99.1 | 99.7 | 99.6 | 71.4 |

Table 3: Hits@1 values of different KGQA methods on MetaQA and WSP. Hop- n denotes the hop- n questions of MetaQA.

MemNN are limited by the error propagation over multi-hop reasoning, i.e., as the number of hops increases, its performance is degraded significantly. GraftNet and PullNet perform similarly well on all datasets (except MetaQA-hop3) as both of them rely on subgraphs to retrieve the answers. Compared to GraftNet, PullNet has much improved results on MetaQA-hop3, indicating that the proposed pull operation is more suitable to complex questions. EmbedKGQA achieves a good performance on MetaQA, but a relatively lower performance on WSP. This is because treating question as a relation path in a triple may introduce more noise especially when the question is more complex. Even though DCRN achieves the best performance on MetaQA-hop2, it still suffers from error propagation when inferring the reasoning paths. For more complex WSP questions, DCRN has a 3.6-point lower accuracy than that of our method. In general, TransferNet is the most competitive one to our SSKGQA. While both methods have the best results on WSP, SKGQA has an improved performance on MetaQA-hop1 over TransferNet, and is almost neck to neck on hop-2 and hop-3. Overall, SSKGQA outperforms or achieves comparable exact-match hits@1 performances to the other methods, demonstrating the effectiveness of our proposed method.

5.5 Ablation Study

We further investigate the effectiveness of different components of SSKGQA, including semantic structure based filtering, Structure-BERT, the BERT-based ranking model, etc.

| | Hop-1 | Hop-2 | Hop-3 | WSP |
|--------|--------------|--------------|--------------|--------------|
| w/o SS | 99.11 | 93.71 | 62.10 | 45.89 |
| w/ SS | 99.26 | 99.03 | 95.69 | 58.51 |

Table 4: Hits@1 values of SSKGQA w/ SS and w/o SS on MetaQA and WSP.

5.5.1 Impact of Semantic Structure based Filtering

One of the core ideas of SSKGQA is the semantic structure based filtering. In this section, we evaluate the effectiveness of this operator by enabling / disabling it and report the final performances of SSKGQA, which correspond to the w/ SS and w/o SS results in Table 4. For the purpose of illustration, when we enable the filtering (w/ SS), we assume that our Structure-BERT classifier can correctly predict the semantic structures of all the questions with a 100% accuracy, and therefore the impact of the filtering isn't affected by the accuracy of the classifier. For ease of experimentation, we use a BiGRU as the ranking model in this experiment.

Table 4 reports the impacts of the semantic structure based filtering. It can be observed that for simple questions, e.g., MetaQA-hop1, SSKGQA w/ SS and w/o SS have very similar performances. However, when the questions are more complex, SSKGQA w/ SS achieves significantly higher accuracies (sometimes over 10%) than SSKGQA w/o SS, demonstrating the effectiveness of the semantic structure based filtering for complex questions.

5.5.2 Accuracy of Structure-BERT

Structure-BERT plays a critical role in SSKGQA as it predicts the semantic structure of a question, which is then used to filter out noisy candidate query graphs. In this section, we evaluate the accuracy of Structure-BERT and compare it with other design choices.

Specifically, we compare the performance of Structure-BERT with four other classifiers, including BiGRU and three pre-trained language models: BERT (Devlin et al., 2018), DistilBERT (Budzianowski and Vulić, 2019) and CamemBERT (Martin et al., 2020)). For these four classifiers, they classify a question directly to one of the six semantic structures without considering topic entity and information fusion as in Structure-BERT.

Table 5 reports the classification accuracies of different classifiers on the questions from MetaQA

| Model | Hop1 | Hop2 | Hop3 | WSP |
|----------------|--------------|--------------|--------------|--------------|
| BiGRU | 96.44 | 94.49 | 98.83 | 80.95 |
| BERT | 94.52 | 98.70 | 96.22 | 82.62 |
| DistilBERT | 95.66 | 98.30 | 97.02 | 83.37 |
| CamemBERT | 96.66 | 97.30 | 98.26 | 81.90 |
| Structure-BERT | 99.24 | 99.87 | 99.73 | 86.97 |

Table 5: Classification accuracies of different classifiers on predicting semantic structures of questions from MetaQA and WSP.

and WSP. As can be seen, our Structure-BERT achieves nearly 100% accuracies on MetaQA and 86.97% accuracy on WSP, demonstrating the effectiveness of Structure-BERT on semantic structure prediction. Further, Structure-BERT consistently outperforms all the other classifiers by notable margins, indicating the importance of leveraging both question and topic entity for information fusion for semantic structure prediction. We also notice that the classification accuracy on WSP is much lower than that of MetaQA. This is likely due to: (1) the class imbalance issue of the WSP questions, and (2) much smaller number of training questions in WSP (3,304) than that of MetaQA (329,282). We will leave the further improvements of Structure-BERT on WSP to future works.

5.5.3 Performance of BERT-based Ranking Model

The BERT-based ranking model decides which candidate query graph is to be used to retrieve the final answer. Therefore, its performance is of the paramount importance to SSKGQA. In this section, we evaluate the effectiveness of our proposed BERT-based ranking model and compare it with other three ranking methods, including 1) CNN (Yih et al., 2015), which uses a CNN to learn the representation of question and candidate query graph for ranking. 2) BiGRU, which uses a BiGRU to learn the representation of question and candidate query graph for ranking. 3) BERT (Devlin et al., 2018), which uses a pre-trained BERT⁵ to extract the representation of question and candidate query graph for ranking.

Table 6 reports the performances of different ranking models on MetaQA and WSP, where BERT* denotes our proposed BERT-based ranking model that has one extra multi-head attention layer on top of the hidden state of the pre-trained

⁵<https://huggingface.co/bert-base-uncased>

| Model | Hop-1 | Hop-2 | Hop-3 | WSP |
|--------------|--------------|--------------|--------------|--------------|
| CNN | 97.70 | 99.21 | 92.91 | 50.24 |
| BiGRU | 98.87 | 98.95 | 95.43 | 56.51 |
| BERT | 99.49 | 99.26 | 99.54 | 71.02 |
| BERT* (ours) | 99.10 | 99.69 | 99.64 | 71.40 |

Table 6: Hits@1 values of different ranking models on MetaQA and WSP. BERT* denotes our BERT-based ranker.

BERT. As we can see, the BERT-based ranking models (BERT and BERT*) outperform the traditional CNN or BiGRU based ranking models since the former can leverage large scale pre-trained BERT for transfer learning. Our BERT* further improves the performance of the pre-trained BERT due to the additional attention layer which enables model to reweight the attention values to different semantic units in the input and enhance the semantic representation of question and candidate query graphs for ranking.

To validate the design choices of our BERT-based model, we run additional ablation studies on different factors of our ranking model, such as number of negative query graphs for the triplet loss based training and number of heads in the added multi-head attention layer. The details are relegated to Sec. 7.3 in Appendix.

6 Conclusions

This paper introduces SSKGQA, a semantic structure based method to predict query graphs from natural language questions. Compared to prior query graph prediction based methods, SSKGQA filters out noisy candidate query graphs based on the semantic structures of input questions. To this end, we define six semantic structures from common questions of the KGQA benchmarks. A novel Structure-BERT classifier is then introduced to predict the semantic structure of each question, and a BERT-based ranking model with a triplet loss is proposed for query graph ranking. Extensive experiments on MetaQA and WSP demonstrate the superior performance of SSKGQA over seven state-of-the-art methods.

As for future work, we plan to investigate techniques to design additional semantic structures to cover the remaining 22.98% of questions in WSP. We would like also to improve Structure-BERT’s accuracy on WSP by addressing the class imbalance and data scarcity issues of WSP.

Acknowledgements

We would like to thank the anonymous reviewers for their comments and suggestions, which helped improve the quality of this paper. We would also gratefully acknowledge the support of VMware Inc. for its university research fund to this research.

References

- Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. Large-scale simple question answering with memory networks. *arXiv preprint arXiv:1506.02075*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcíaaduran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *International Conference on Neural Information Processing Systems*.
- Paweł Budzianowski and Ivan Vulić. 2019. Hello, it’s gpt-2—how can i help you? towards the use of pre-trained language models for task-oriented dialogue systems. *arXiv preprint arXiv:1907.05774*.
- Jianyu Cai, Zhanqiu Zhang, Feng Wu, and Jie Wang. 2021. Deep cognitive reasoning network for multi-hop question answering over knowledge graphs. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 219–229.
- Yongrui Chen, Huiying Li, Yuncheng Hua, and Guilin Qi. 2020. Formal query building with query structure prediction for complex question answering over knowledge base. In *International Joint Conference on Artificial Intelligence (IJCAI)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- David Golub and Xiaodong He. 2016. Character-level question answering with attention. *arXiv preprint arXiv:1604.00727*.
- Aisha Urooj Khan, Amir Mazaheri, Niels Da Vitoria Lobo, and Mubarak Shah. 2020. Mmft-bert: Multimodal fusion transformer with bert encodings for visual question answering. *arXiv preprint arXiv:2010.14095*.
- Nabila A Khodeir. 2021. Bi-gru urgent classification for mooc discussion forums based on bert. *IEEE Access*, 9:58243–58255.
- Yunshi Lan and Jing Jiang. 2020. Query graph generation for answering multi-hop complex questions from knowledge bases. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Gaurav Maheshwari, Priyansh Trivedi, Denis Lukovnikov, Nilesh Chakraborty, Asja Fischer, and Jens Lehmann. 2019. Learning to rank query graphs for complex question answering over knowledge graphs. In *International semantic web conference*, pages 487–504. Springer.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villenote de la Clergerie, Djamé Seddah, and Benoît Sagot. 2020. Camembert: a tasty french language model. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Alexander Miller, Adam Fisch, Jesse Dodge, Amir-Hossein Karimi, Antoine Bordes, and Jason Weston. 2016. Key-value memory networks for directly reading documents. *arXiv preprint arXiv:1606.03126*.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. 2009. Semantics and complexity of sparql. *ACM Transactions on Database Systems (TODS)*, 34(3):1–45.
- Kechen Qin, Cheng Li, Virgil Pavlu, and Javed Aslam. 2021. Improving query graph generation for complex question answering over knowledge base. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4201–4207.
- Nils Reimers and Iryna Gurevych. 2019. Sentencebert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Apoorv Saxena, Aditay Tripathi, and Partha Talukdar. 2020. Improving multi-hop question answering over knowledge graphs using knowledge base embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4498–4507.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.
- Jiaxin Shi, Shulin Cao, Lei Hou, Juanzi Li, and Hanwang Zhang. 2021. Transfernnet: An effective and transparent framework for multi-hop question answering over relation graph. *arXiv preprint arXiv:2104.07302*.
- Saurabh Srivastava, Mayur Patidar, Sudip Chowdhury, Puneet Agarwal, Indrajit Bhattacharya, and Gautam Shroff. 2021. Complex question answering on knowledge graphs using machine translation and

multi-task learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3428–3439, Online. Association for Computational Linguistics.

Haitian Sun, Tania Bedrax-Weiss, and William W Cohen. 2019a. Pullnet: Open domain question answering with iterative retrieval on knowledge bases and text. *arXiv preprint arXiv:1904.09537*.

Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Kathryn Mazaitis, Ruslan Salakhutdinov, and William W Cohen. 2018. Open domain question answering using early fusion of knowledge bases and text. *arXiv preprint arXiv:1809.00782*.

Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019b. Rotate: Knowledge graph embedding by relational rotation in complex space. *arXiv preprint arXiv:1902.10197*.

Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *International Conference on Machine Learning*, pages 2071–2080. PMLR.

Kun Xu, Yuxuan Lai, Yansong Feng, and Zhiguo Wang. 2019. Enhancing key-value memory neural networks for knowledge based question answering. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.

Naganand Yadati, RS Dayanidhi, S Vaishnavi, KM Indira, and G Srinidhi. 2021. Knowledge base question answering through recursive hypergraphs. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 448–454.

Scott Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. 2015. Semantic parsing via staged query graph generation: Question answering with knowledge base.

Wen Tau Yih, Matthew Richardson, Chris Meek, Ming Wei Chang, and Jina Suh. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*.

Yuyu Zhang, Hanjun Dai, Zornitsa Kozareva, Alexander J Smola, and Le Song. 2018. Variational reasoning for question answering with knowledge graph. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

7 Appendix

7.1 Example Questions and their Semantic Structures

We provide six example questions and their corresponding semantic structures. These examples are selected from the MetaQA and WSP benchmarks. E denotes a topic entity, r refers to all types of relations, v is an existential variable, and C denotes a constraint.




Figure 6: Examples questions and their semantic structures

7.2 Semantic Structure Annotation for Training Questions

We need to annotate each training question to a semantic structure in order to train Structure-BERT for semantic structure prediction. Here we describe how we can automatically annotate each training question in MetaQA and WSP.

7.2.1 MetaQA

Annotating semantic structure for each training question in MetaQA is straightforward. Since training questions are organized by number of hops, all training questions with 1-hop are labeled to SS1, the 2-hop ones to SS2, and the 3-hop ones to SS3, respectively.

7.2.2 WSP

The SPARQL command for each question is provided in WSP. Thus, we can readily extract the query graph of a question from its SPARQL command. See an example in Figure 7, where a query graph is shown on the left and the corresponding

SPARQL command is shown on the right; the correspondences between two parts are marked by red lines. Once the query graph is extracted, each training question can be readily annotated to a semantic structure based on its query graph.



Figure 7: Mapping from a SPARQL command to its query graph

7.3 Additional Ablation Study

| | | | | |
|----------------|--------------|-------|--------------|-------|
| Number of Neg. | 1 | 5 | 10 | 50 |
| hits@1 | 64.57 | 68.87 | 70.08 | 70.79 |
| Number of Neg. | 100 | 200 | 300 | 500 |
| hits@1 | 71.29 | 71.23 | 71.29 | 70.58 |

Table 7: Hits@1 values on WSP of our BERT-based ranking model when trained with different number of negative query graphs.

Number of negative query graphs Given a question, a plenty of negative query graphs can be generated by enumeration from a KG. By analyzing the questions in WSP, we found that the maximal number of negative query graphs that can be extracted for WSP is around 500. We need to determine a proper number of negative samples to train the ranking model with the triplet loss. To this end, we evaluate the performances of our BERT-based ranker when trained with different numbers of negative samples. The results are reported in Table 7, where 8 different number of negative samples are considered. As we can see, when number of negative samples is over 100, the ranking model achieves improved hits@1 performances; when $n = 100$ our BERT-based ranking model yields the best hits@1 with a good run-time performance.

| | | | |
|-----------------|-------|--------------|-------|
| Number of Heads | 1 | 3 | 6 |
| hits@1 | 71.29 | 71.40 | 70.63 |

Table 8: Hits@1 values on WSP of our BERT-based ranking model with different number of attention heads in the added attention layer.

Number of heads Another design choice for the extra multi-head attention layer in our BERT-based

ranker is the number of attention heads. We therefore evaluate the performance of our BERT-based ranker with 3 different number of attention heads. The results in Table 8 shows that when the number of attention heads is 3, our BERT-based ranker achieves the best hits@1.