

# Knowledge Graph Question Answering using Graph-Pattern Isomorphism

Daniel VOLLMERS<sup>a</sup> Rricha JALOTA<sup>a</sup> Diego MOUSSALLEM<sup>a,d</sup>  
Hardik TOPIWALA<sup>a</sup> Axel-Cyrille NGONGA NGOMO<sup>a</sup> and Ricardo USBECK<sup>b,c</sup>

<sup>a</sup>*Data Science Group, Paderborn University, Germany*

<sup>b</sup>*Fraunhofer IAIS, Dresden, Germany*

<sup>c</sup>*Universität Hamburg, Germany*

<sup>d</sup>*Globo, Rio de Janeiro, Brazil*

**Abstract.** Knowledge Graph Question Answering (KGQA) systems are often based on machine learning algorithms, requiring thousands of question-answer pairs as training examples or natural language processing pipelines that need module fine-tuning. In this paper, we present a novel QA approach, dubbed TeBaQA. Our approach learns to answer questions based on graph isomorphisms from basic graph patterns of SPARQL queries. Learning basic graph patterns is efficient due to the small number of possible patterns. This novel paradigm reduces the amount of training data necessary to achieve state-of-the-art performance. TeBaQA also speeds up the domain adaption process by transforming the QA system development task into a much smaller and easier data compilation task. In our evaluation, TeBaQA achieves state-of-the-art performance on QALD-8 and delivers comparable results on QALD-9 and LC-QuAD v1. Additionally, we performed a fine-grained evaluation on complex queries that deal with aggregation and superlative questions as well as an ablation study, highlighting future research challenges.

**Keywords.** Question Answering, Basic Graph Pattern, Isomorphism, QALD

## 1. Introduction

The goal of most Knowledge Graph (KG) Question Answering (QA) systems is to map natural language questions to corresponding SPARQL queries. This process is known as semantic parsing [5] and can be implemented in various ways. A common approach is to utilize query templates (alias graph patterns) with placeholders for relations and entities. The placeholders are then filled with entities and relations extracted from a given natural language question [1,16,26] to generate a SPARQL query, which is finally executed. Semantic parsing assumes that a template can be constructed or chosen to represent a natural language question’s internal structure. Thus, the KGQA task can be reduced to finding a matching template and filling it with entities and relations extracted from the question.

The performance of KGQA systems based on this approach depends heavily on the implemented query templates, which depend on the question’s complexity and the KG’s topology. Consequently, costly hand-crafted templates designed for a particular KG cannot be easily adapted to a new domain.

In this work, we present a novel KGQA engine, dubbed TeBaQA. TeBaQA alleviates manual template generation’s effort by implementing an approach that learns templates from existing KGQA benchmarks. We rely on learning of templates based on isomorphic basic graph patterns.

The goal of TeBaQA is to employ machine learning and feature engineering to learn to classify natural language questions into isomorphic basic graph pattern classes. At execution time, TeBaQA uses this classification to map a question to a basic graph pattern, i.e., a template, which it can fill and augment with semantic information to create the correct SPARQL query.

TeBaQA achieves state-of-the-art performance partially without any manual effort. In contrast to existing solutions, TeBaQA can be easily ported to a new domain using only benchmark datasets, partially proven by our evaluation over different KGs and train-test-sets. We use a best-effort to work with the data at hand instead of either (i) requiring a resource-intensive dataset creation and annotation process to train deep neural networks or (ii) to hand-craft mapping rules for a particular domain. Our contributions can be summarized as follows:

- We present TeBaQA, a QA engine that learns templates from benchmarks based on isomorphic basic graph patterns.
- We describe a greedy yet effective ranking approach for query templates, aiming to detect the best matching template for a given input query.
- We evaluate TeBaQA on several standard KGQA benchmark datasets and unveil choke points and future research directions.

The code is available at <https://github.com/dice-group/TeBaQA> and a demo of TeBaQA over encyclopedic data can be found at <https://tebaqa.demos.dice-research.org/>. We also provide an online appendix which contains more details about our algorithms and their evaluations at [https://github.com/dice-group/TeBaQA/blob/master/TeBaQA\\_appendix.pdf](https://github.com/dice-group/TeBaQA/blob/master/TeBaQA_appendix.pdf).

## 2. Related Work

The domain of Knowledge Graph Question Answering has gained traction over the last few years. There has been a shift from simple rule-based systems to systems with more complex architectures that can answer questions with varying degrees of complexity. In this section, we provide an overview of the recent work. We begin with approaches that took part in QALD challenge series [29,28,27].

gAnswer2 [35] addresses the translation of natural language to SPARQL as a sub-graph matching problem. Following the rule-based paradigm, QAnswer [9] utilizes the semantics embedded in the underlying KG (DBpedia and Wikidata) and employs a combinatorial approach to create SPARQL queries from natural language questions. QAnswer overgenerates possible SPARQL queries and has to learn the correct SPARQL query ranking from large training datasets.

Second, we introduce approaches that rely on templates to answer natural questions over knowledge bases. Hao et al. [15] introduced a pattern-revising QA system for answering simple questions. It relies on pattern extraction and joint fact selection, enhanced by relation detection, for ranking the candidate subject-relation pairs. NEQA [1] is a

template-based KBQA system like TeBaQA, which uses a continuous learning paradigm to answer questions from unseen domains. Apart from using a similarity-based approach for template-matching, it relies on user feedback to improve over time. QUINT [2] also suggests a template learning system that can perform on compositional questions by learning sub-templates. KBQA [7] extracted 27 million templates for 2782 intents and their mappings to KG relations from a QA corpora (Yahoo! Answers). The way these templates are generated and employed differs from that of TeBaQA. NEQA relies on active learning on sub-parts and thereby, possibly misses the semantic connection between question parts. KBQA learns many templates and can hence fail to generalize well to other domains or KB structures.

In 2019, Zheng et al. [34] use structural query pattern, a coarse granular version of SPARQL basic graph patterns which they later augment with different query forms and thus can generate also SPARQL query modifiers. The closest work to ours is by Athreya et al. [3] based on a tree-based RNN to learn different templates on LC-QuAD v1 which the authors directly derive from the LC-QuAD v1 inherent SPARQL templates and thus cannot generalize to other KGs or datasets. CONQUEST [4] is an enterprise KGQA system which also assumes the SPARQL templates are given. It then matches the questions and templates by vectorizing both and training one classifier, namely Gaussian Naïve Bayes.

QAMP [31] is an unsupervised message passing system using a similarly simple approach to information extraction as TeBaQA. QAMP outperforms QAnswer on LC-QuAD v1. Recently, Kapanipathi et al. [17] present a system, dubbed NSQA, which is in pre-print. NSQA is based on a combination of semantic parsing and reasoning. Their modular approach outperforms gAnswer and QAnswer on QALD-9 as well as QAnswer and QAMP on LC-QuAD V1. In contrast to seminal semantic parsing work, e.g., by Berant et al. [5], we assume relatively small training data; thus, learning a generalization via semantic building blocks will not work and consequently learn the whole template on purpose.

Third, there are also other QA approaches which work with neural networks relying on large, templated datasets such as sequence-to-sequence models [24,33]. However, we do not focus on this research direction in this work. We refer the interested reader to extensive surveys and overview papers such as [6,10,16].

### 3. Approach

TeBaQA is based on isomorphic graph patterns, which can be extracted across different SPARQL queries and hence be used as templates for our KGQA approach. Figure 1 provides an overview of TeBaQA's architecture and its five main stages:

First, all questions run through a **Preprocessing** stage to remove semantically irrelevant words and create a set of meaningful n-grams. The **Graph-Isomorphism Detection and Template Classification** phase uses the training sets to train a classifier based on a natural language question and a SPARQL query by analyzing the basic graph pattern for graph isomorphisms. The main idea is that structurally identical SPARQL queries represent syntactically similar questions. At runtime, a question is classified into a ranked list of SPARQL templates. While **Information Extraction**, TeBaQA extracts all critical information such as entities, relations and classes from the question and determines the

answer type based on a KG-agnostic set of indexes. In the **Query Building** phase, the extracted information are inserted into the top templates, the SPARQL query type is determined, and query modifiers are added. The resulting SPARQL queries are executed, and their answers are compared with the expected answer type. The subsequent **Ranking** is based on a combination of all information, the natural language question and the returned answers.

In the following, we present each of these steps in more detail. We use DBpedia [18] as reference KG for the sake of simplicity in our description.

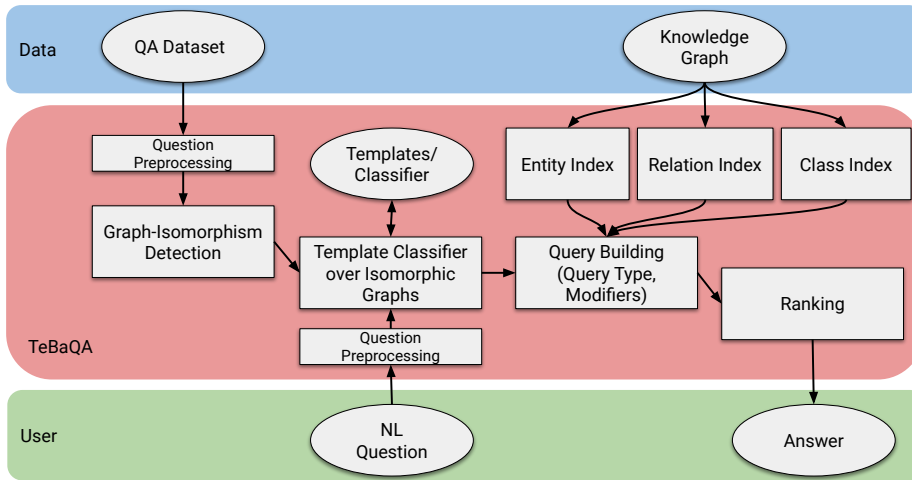


Figure 1. TeBaQA architecture on the running example.

### 3.1. Question Preprocessing

There are often words that do not contribute any information to the answer to natural language questions. Thus, we distinguish semantically relevant and irrelevant n-grams. Irrelevant n-grams can lead to errors that could propagate through the architecture. An example of this is the entity `dbr:The_The`<sup>1</sup>. If the word *The* were to be wrongly associated with this entity every time *the* occurs in a question, the system’s performance would decrease severely. However, irrelevant words are sometimes part of entities, e.g., `dbr:The_Two_Towers`, so we cannot always filter these words. For this reason, we combine up to six neighboring words from the question to n-grams and remove all n-grams that contain stop words only. To identify irrelevant words, we provide a stop word list that contains the most common words of a particular language that are highly unlikely to add semantic value to the sentence. Additionally, TeBaQA distinguishes relevant and irrelevant n-grams using part-of-speech (POS) tags. Only n-grams beginning with JJ, NN or VB POS-tags are considered as relevant. After this preprocessing step, TeBaQA maps the remaining n-grams to entities from DBpedia in the information extraction step.

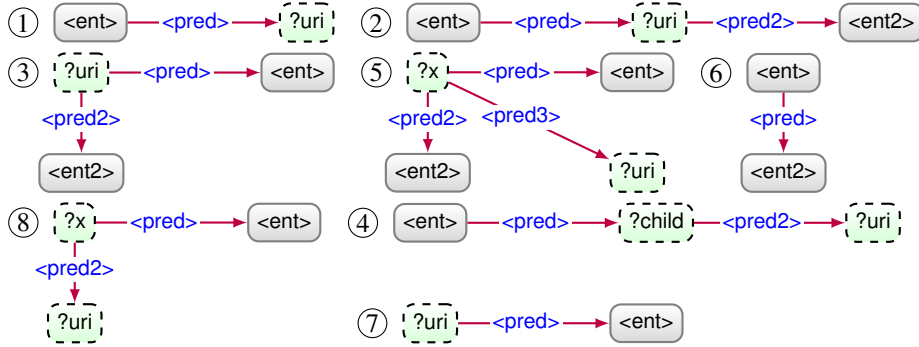
<sup>1</sup>`dbr:` is a prefix which stands for `http://dbpedia.org/resource/`

### 3.2. Graph-Isomorphism Detection and Template Classification

TeBaQA classifies a question to determine each isomorphic basic graph pattern (BGP). Since SPARQL is a graph-based query language [21], the structural equality of two SPARQL queries can be determined using an isomorphism. At runtime, TeBaQA can classify incoming questions to find the correct query templates, in which later semantic information can be inserted at runtime.

#### 3.2.1. SPARQL BGP Isomorphism to Create Template Classes:

Using the training datasets, TeBaQA generates one basic graph pattern for each given question and its corresponding SPARQL query, see Figure 2. Subsequently, all *isomorphic* SPARQL queries are grouped into the same class. Now, each class contains semantically different natural language questions but structurally similar SPARQL queries.



**Figure 2.** All basic graph patterns used as classes for QALD-8 and QALD-9 which later become templates. Note, the depicted templates contain more than five examples in the training dataset. Our running example, "Who was the doctoral advisor of Albert Einstein?" belongs to template (1).

**Theorem 1 (Isomorphism of labeled graphs)** *Two labeled graphs are isomorphic when a 1:1 relationship and surjective function is present between the nodes of the graphs, wherein the node labels, edge labels, and neighborhood relationships are preserved by the mapping [13].*

#### 3.2.2. Question Features and Classification

Next, TeBaQA trains a classifier that uses all questions of an isomorphism class as input to calculate features for this class. A feature vector holds all the information required to make a reliable statement about which question belongs to which class. The features can be seen in Table 1.

The features can be divided into semantic and syntactic features. *QuestionWord*, *EntityPerson* and *QueryResourceType* form the group of semantic features and represent particular content aspects of the question, e.g., persons or specific topics that are mentioned in the question. All other features describe the structure of the question.

Note, other features were investigated, which did not improve the model's recognition rate. We report these features to aid future research in this area: 1) Cultural Categories: Mainly included music and movies, e.g., *Who is the singer on the album The*

**Table 1.** Features to map a question to an isomorphic basic graph pattern.

Feature	Type	Description
QuestionWord	Nominal	Adds the question word (e.g. Who, What, Give) as a feature.
EntityPerson	Boolean	Checks the named entity tags of the sentence to see if any persons are mentioned in it.
NumberOfToken	Numeric	Stores the number of tokens separated by spaces excluding punctuation.
QueryResourceType	Nominal	Categorizes the question based on a list of subject areas, e.g., film, music, book or city.
Noun	Numeric	Aggregates the number of nouns.
Number	Numeric	Indicates how often numbers occur in the question.
Verb	Numeric	Aggregates the number of verbs.
Adjective	Numeric	Aggregates the number of adjectives.
Comperative	Boolean	Indicates whether comparative adjectives or adverbs are included in the sentence.
TripleCandidates	Numerical	Estimates how many SPARQL triples are needed to answer the question based on the number of verbs, adjectives, and related nouns.

*Dark Side of the Moon?* and 2) Geographical entities: Questions in which countries or cities occur, as well as where questions, e.g., *In which country is Mecca located?*

Using features above, it is possible to represent the question *Who was the doctoral advisor of Albert Einstein?* with the following vector:

<Who, Person , 8 , dbo : Person , 1 , 0 , 1 , 0 , NoComperative , 1 >
---

TeBaQA trains a statistical classifier using the described features extracted from the input question and the isomorphic basic graph patterns as class labels. A feature vector’s target class can be determined by generating the basic graph pattern for the corresponding SPARQL query and assigning the class, which represents this pattern. An evaluation can be found in Section 4.2.

### 3.3. Information Extraction

TeBaQA identifies entities, classes, and relations to fill the placeholders of a particular SPARQL template. Since questions are shorter than usual texts, semantic entity linking tools such as DBpedia Spotlight [8] or MAG [20,19] do not perform correctly due to the lack of semantic context information. For example, in *Who was the doctoral advisor of Albert Einstein?*, the word *Einstein* has to be linked to `dbr:Albert_Einstein` and not to any other person with that name. For this reason, we apply a KB-agnostic search index-based approach to identify entity, relation and class candidates. TeBaQA uses three indexes, which are created before runtime.

#### 3.3.1. Entity Index

The entity index contains all entities from the target knowledge graph. To map an n-gram from the preprocessing step to an entity, TeBaQA queries against the index’s label field. The index contains information about entities, relations and classes connected to the entity at hand.

September 2021

### 3.3.2. Relation Index and Class Index

These two indexes contain all OWL classes and relations from a KG. The indexes are used to map n-grams to relations and classes of the ontologies of the KB. TeBaQA additionally indexes hypernyms and synonyms for all relations and classes.<sup>2</sup>

Consider the question *Who was the doctoral mentor of Einstein?*. DBpedia contains only the relation `dbo:doctoralAdvisor`<sup>3</sup> and not `dbp:mentor`<sup>4</sup>. Through the synonym *advisor* of *mentor*, the relation `dbo:doctoralAdvisor` can be determined. This example highlights the lexical and semantic gap between natural language and knowledge graphs.

### 3.3.3. Disambiguation

By querying the indexes for an n-gram, we get candidates for entities, relations and classes, whose labels contain all tokens of the n-gram. Since a candidate's label may contain more tokens than the n-gram, we apply a Levenshtein-distance filter of 0.8 on the candidates. All remaining candidates are used to fill a given template.

## 3.4. Query Building

### 3.4.1. Template Filling

For filling the templates, we facilitate the information about connected entities and connected relations for the found entities from the entity-index. For the triples in a template, there are two cases:

1.) The triple contains one placeholder for an entity and one placeholder for a relation. In this case, we resort to only the connected relation information from the entity index. An entity candidate  $e$  and a relation candidate  $p$  are combined to a triple  $\langle e, p, ?v \rangle$  or  $\langle ?v, p, e \rangle$  if the set of connected relations  $S(e)$  of the entity  $e$  contains  $p$  and if the connected n-grams do not contain each other.

2.) The triple contains only one placeholder for a relation  $p'$ . This case only occurs when at least one triple in the template matches case 1. We can utilize these triples to generate matching triples for the given triple. Thus, we query the entity index and search for a set of entities  $S(e')$  connected with the entity  $e$  by the relation  $p$ . All connected relations from the entities in  $S(e')$  in the set of relation candidates and whose n-grams do not cover the n-grams of  $e$  and  $p$  are candidates for  $p'$ .

Each candidate SPARQL query is checked for consistency with the ontology. In general, there are query patterns that do not contain variables. This case only occurs in ask queries like *Did Socrates influence Aristotle?*. We ignore this case to keep simplicity and aware of the performance impact. To summarize, TeBaQA creates several candidate SPARQL queries per template and thus per question.

---

<sup>2</sup>The dictionary can be found at <https://github.com/dice-group/NLIWOD/tree/master/qa.annotation/src/main/resources> which was previously used by [23]

<sup>3</sup>`dbo:` stands for <http://dbpedia.org/ontology/>

<sup>4</sup>`dbp:` stands for <http://dbpedia.org/property/>

### 3.4.2. Query Modifiers and Query Types

To translate a question into a semantically equivalent SPARQL query, it is often not enough to recognize the entities, relations or classes in the question and insert them into a SPARQL query. For example, the question *How many children did Benjamin Franklin have?* asks for the number of children and not the concrete list. Thus, we apply a rule-based look-up to add query modifiers and choose a query type. The supported modifiers are *COUNT*, if the question contains keywords like *How many or How much*, *Filter(?x <?y)*, if we identify comparatives and *ORDERBY [ASC(?x) | DESC(?x)] LIMIT 1*, if the question contains superlatives. Additionally, we support ASK-type questions, if keywords like *Is* or *Are* are identified.

Note, the templates and their respective basic graph pattern neither contain information about the query type nor information about query modifiers. Thus, TeBaQA generates one SPARQL query per candidate SPARQL query for each cross-product of query type and modifier that is recognized. The outcome is a list of executable queries for each input question.

### 3.5. Ranking

Since the *conciseness* of an answer plays a decisive role in QA in contrast to full-text search engines, only the answer that corresponds best to the user's intention should be returned. Thus, all generated SPARQL queries and their corresponding answers are ranked. This ranking is carried out in two steps. First, we filter by 1) the expected answer type of the question in comparison to the actual answer type of the query and by 2) the cardinality of the result set. Second, TeBaQA ranks the quality of the remaining SPARQL queries.

#### 3.5.1. Answer Type and Cardinality Check:

For certain types of answer sets, only those that match the question's expected answer type are considered for the next ranking step. We empirically analyzed the benchmark datasets and derived a rule-based system for the most common expected answer type and their distinguishing features.

- Temporal questions usually begin with the question word *When*, e.g., *When was the Battle of Gettysburg?*. TeBaQA expects a date as the answer type.
- Decision questions mostly start with a form of *Be*, *Do* or *Have*. The possible answer type is boolean.
- Questions that begin with *How much* or *How many* can be answered with numbers. This also includes questions that begin with a combination of the question word *How* and a subsequent adjective such as *How large is the Empire State Building?*

If none of the above rules apply to a question, the result set's cardinality is checked. There are two cases: First, when several answers are needed to answer a question fully. Consider *Which ingredients do I need for carrot cake?*, if only one answer is found for this question, it can be assumed that it is either wrong or incomplete. Second, when there is only one answer to a question, e.g., *In which UK city are the headquarters of the MI6?*, an answer consisting of several entities would not be correct.



To recognize to which query type (ASK or SELECT) a question belongs, the first noun or the first compound noun after the question word is checked. If they occur in the singular form, a single answer is needed to answer the question. For the above question *In which UK city are the headquarters of the MI6?*, the compound noun would be *UK city*. Since both words occur in the singular, it is assumed that only a single answer is required. If the first noun or group of nouns occurs in the plural, this indicates that the question requires multiple answers. In the question *Which ingredients do I need for carrot cake?* the decisive word *ingredients* is in the plural.

However, the answer type of question may not be recognized correctly. For instance, if the question is grammatically correct but words whose singular form is identical to the plural form exist, such as *news*, we cannot determine the correct answer type. These issues will be tackled in future research.

Once the type of question and answer has been determined, all the answers whose type or cardinality does not match the question will be discarded.

### 3.5.2. Quality Ranking:

For the remaining SPARQL queries, TeBaQA calculates a *rating* based on the sum of the individual scores of the bindings  $B$  and the input question *phrase*. A binding  $B$  is the mapping of entities, relations and classes to placeholders contained in one SPARQL query  $q$ . To compute the relatedness factor  $r$ , the following factors are taken into account:

- **Annotation Density:** The annotation density measures that the more words from the sentence are linked to an entity, class or relation, the more likely it is that it corresponds to the intention of the user. For the question *What is the alma mater of the chancellor of Germany Angela Merkel?* one candidate query may apply the binding `dbr:Angela`, while another query applies the binding `dbr:Angela_Merkel`. The former refers only to the word *Angela*. The latter refers to two words of the sentence: *Angela Merkel* and covers longer parts of the *phrase*.
- **Syntactic Similarity:** The syntactic similarity is an indicator of how similar a n-gram of the sentence and the associated binding are. For example, in the question *Who is the author of the interpretation of dreams?* the n-gram *the interpretation of dreams* can be linked with `dbr:The_Interpretation_of_Dreams` or `dbr:Great_Book_of_Interpretation_of_Dreams` among others. The former has a smaller Levenshtein distance and a greater syntactic similarity with the selected n-gram.

We cover both aspects with the following formulas:

$$rating = \sum_{B \in q} r(B, phrase)$$

$$r(entity, phrase) = |words(phrase)| - levenshtein\_ratio(label(B), phrase)$$

After all entities, classes, and relations used in a query have been evaluated and summed up, the rating is corrected down by 30% if more than 50 results are returned by the query, based on empirical observations in the datasets.

## 4. Evaluation

### 4.1. Datasets

We performed the evaluation on the 8th and 9th Question Answering over Linked Data challenge training datasets (QALD-8 train [28] and QALD-9 train [27]), which contain 220 (QALD-8) and 408 (QALD-9) heterogeneous training questions. Additionally, we evaluated on the two LC-QuAD [25,11] datasets with 4000 train and 1000 test questions and 24.000 train and 6.000 test questions, respectively. Across datasets, the questions vary in complexity since they also include comparatives, superlatives, and temporal aggregations. An example of a simple question is *How tall is Amazon Eve?*. A more complex example is *How many companies were founded in the same year as Google?* since it contains a temporal aggregation (*the same year*). We created separate instances of TeBaQA for each of the training datasets and evaluated each instance on the corresponding test dataset.

### 4.2. Classification Evaluation

For the QALD-8 and QALD-9 datasets, eight classes were identified for each dataset, compare Figure 2. For LC-QuAD v1 and LC-QuAD v2, TeBaQA identified 17 classes for both datasets. Since the two LC-QuAD datasets were constructed for more diversity, the classification is more challenging. Note, we omitted classes with less than five examples in the training dataset. We are aware that we are trading our overall performance for classification accuracy.

To this end, we evaluated a variety of machine learning methods, which required questions to be converted into feature vectors. In particular, we used the QALD-8 and QALD-9 training datasets as our training data and used 10-fold cross-validation to evaluate the computed models. All duplicate questions and questions without SPARQL queries were removed from the training datasets. We tested multiple machine learning algorithms with the WEKA framework<sup>5</sup> [12] on our training data using 10-fold cross-validation. To achieve comparable results, TeBaQA uses only the standard configuration of the algorithms. The macro-weighted F-Measure for one fold in cross-validation is calculated from the classes' F-Measures, weighted according to the size of the class. After that, we calculated the average macro-weighted F-Measure across all folds. On the QALD datasets, the algorithm *RandomizableFilteredClassifier* achieves the highest F-Measures of 0.523964 (QALD-8) and 0.528875 (QALD-9), respectively. On LC-QuAD v1, we achieve a template classification f-measure of 0.400464 and 0.425953 on LC-QuAD v2 using a *MultilayerPerceptron*. Consequently, we use the best performing classifiers for the end-to-end evaluation.

Similar experiments can be found in Athreya et al.'s work [3]. The authors use a recurrent neural network, i.e., a tree-LSTM, to identify templates in LC-QuAD and achieve an accuracy of 0.828 after merging several template classes manually.

---

<sup>5</sup><https://www.cs.waikato.ac.nz/ml/weka/>

### 4.3. GERBIL QA Benchmark

For evaluation, we used the FAIR benchmarking platform GERBIL QA [30] to ensure future reproducibility of the experiments.

Table 2 contains the results of selected Question Answering systems, measured against the QALD-8 and the QALD-9 test benchmarks<sup>6</sup>. We focused on English questions only, as English is supported by all available QA systems at the time of these experiments. The macro value for Precision, Recall and F-Measure was selected. The evaluation was performed with GERBIL version 0.2.3 if possible. We report always the highest numbers if several papers reported numbers and evaluation with GERBIL was not possible.

**Table 2.** Results of TeBaQA and other state-of-the-art QA systems for multiple datasets. \* indicates F-1 measure instead of QALD F-Measure [30]. \*\* Numbers taken from the corresponding paper.

System	KB	Dataset	Precision	Recall	QALD F-Measure	Avg. Time in s
gAnswer2 [35]	DBpedia	QALD-8	0.337	0.354	0.440	4.548
QAnswer [9]	DBpedia	QALD-8	0.452	0.480	0.512	0.446
Zheng et al. [34]**	DBpedia	QALD-8	0.459	0.463	* 0.461	-
TeBaQA	DBpedia	QALD-8	0.476	0.488	<b>0.556</b>	28.990
Elon [27]	DBpedia	QALD-9	0.049	0.053	0.100	0.219
gAnswer [35]	DBpedia	QALD-9	0.293	0.327	0.430	3.076
NSQA [17]**	DBpedia	QALD-9	0.314	0.322	*0.453	-
QAnswer [9]	DBpedia	QALD-9	0.261	0.267	0.289	0.661
QASystem [27]	DBpedia	QALD-9	0.097	0.116	0.200	1.014
Zheng et al. [34]**	DBpedia	QALD-9	0.458	0.471	<b>*0.463</b>	-
TeBaQA	DBpedia	QALD-9	0.241	0.245	0.374	5.049
NSQA [17]**	DBpedia	LC-QuAD v1	0.382	0.404	*0.383	-
QAMP [31]**	DBpedia	LC-QuAD v1	0.250	0.500	*0.330	0.720
QAnswer [9] **	DBpedia	LC-QuAD v1	0.590	0.380	<b>*0.460</b>	1.500
TeBaQA	DBpedia	LC-QuAD v1	0.230	0.229	0.300	36.000
TeBaQA	Wikidata	LC-QuAD v2	0.140	0.136	0.227	-

On the QALD-8 benchmark, TeBaQA achieved the best results in terms of F-Measure by 5% QALD F-measure. Our average time is significantly larger than the other reported systems since the ranking mechanism gets activated and then fires several SPARQL queries after the initial null-retrieving SPARQL query.

On QALD-9, TeBaQA is in fourth place with a QALD F-Measure of 0.37. This implies that TeBaQA achieves comparable or partially better results than other semantic QA systems with a wide margin of possible improvements, as shown in the ablation study. QALD-9 is a more challenging benchmark than QALD-8 since it contains many questions that require complex queries with more than two triples. A more in-depth analysis shows that questions from QALD-9 often require complex templates that are not contained in the training queries or have only low support in the training set [14]. This mismatch leads to a high number of misclassifications and explains the limited perfor-

<sup>6</sup>The links to our GERBIL-experiments can be found on our Github page: <https://github.com/dice-group/TeBaQA/blob/master/README.md>

mance on QALD-9 compared to QALD-8 and, thus, its limited generalization abilities to unseen templates. Although we are not outperforming the state-of-the-art systems on QALD-9, TeBaQA a novel research avenue w.r.t. learning from data.

On the LC-QuAD dataset, which contains the most complex questions, TeBaQA achieves an F-Measure of 0.30. We ran this benchmark only once in our system and had some errors during runtime. We will further investigate the performance on LC-QuAD in our future research. Note, we ran LC-QuAD only once through the system to test our independence of the dataset similar to the methodology of Berant et al. [32].

To the best of our knowledge and after reaching out to the authors of LC-QuAD, we are the first system to evaluate on LC-QuAD v2. Since LC-QuAD v2 uses Wikidata instead of DBpedia, we generated separate indexes for the Wikidata KG. The rest of the system remained unchanged. This shows, that TeBaQA can be easily adapted to other Knowledge Graphs.

#### 4.4. SPARQL operation modifiers

The Achilles heel of most QA systems is their inability to deal with SPARQL operation modifiers, e.g., questions involving aggregations or superlatives [22]. In contrast to those other systems, TeBaQA contains functionalities to identify the query type and other modifiers. We analyzed the results of TeBaQA on these benchmark questions and found that TeBaQA was able to answer many of those questions, while other systems like QAnswer and gAnswer fail on this; see supplementary material.

#### 4.5. Ablation Study

**Table 3.** Ablation study for TeBaQA on the QALD-9 test benchmark.

QA System	Precision	Recall	Avg. Time	QALD F-Measure
Perfect classification	0.205	0.210	4.618	0.337
Perfect classification + EL	0.407	0.407	0.355	0.578
Perfect classification + ranking	0.245	0.257	4.029	0.399
Perfect EL	0.301	0.317	0.713	0.473
Perfect EL + ranking	0.302	0.320	0.653	0.477
Perfect ranking	0.258	0.270	6.281	0.405
Perfect classification + EL + ranking	0.407	0.407	0.251	0.578

We performed an ablation study to find the modules of TeBaQA’s pipeline, which influence the end-to-end performance the most. Since the number of possible entities is roughly a magnitude larger than the number of relations and for the sake of experiment time, we omitted to test for perfect relation and class linking.

For the perfect classification experiment, the QALD F-Measure is lower than for the overall system, see Table 2. Investigating the detailed outputs, TeBaQA selects the simple templates containing only one triple more often than the more complex templates because they have more instances, i.e., support, on the QALD-9 train dataset. In many cases, a simple query generates a result set that is a superset of the target result set and, in consequence, decreases the precision.

When TeBaQA fails to fill the correct complex template with the correct entities, the query result set is often disjoint from the target result set. It is also reasonable that TeBaQA fails to fill the more complex templates due to missing or incorrect entity links.

Still, there is a gap in the system, which becomes evident when looking at the perfect classification plus ranking. Ranking gets only activated if the perfect template filled with semantic information does not retrieve any answer. That is, TeBaQA fails to find the correct modifiers or needs to circle through other semantic information candidates.

When the perfect classification is combined with perfect entity linking, the results reach a QALD F-measure of 0.58, which clearly would outperform any other system. The same happens if we add the perfect ranking. The results are the same in both cases because the ranking is most often not triggered since the perfect template is already filled correctly.

The strongest, single influence is the entity linking part, enabling TeBaQA to jump to 0.47 F-measure. We will tackle this challenging module in future work.

Regarding runtime, failing to find the perfect template and then iterating through the ranking, i.e., querying the SPARQL endpoint often, increases the average time needed significantly.

## 5. Summary and Future Work

We presented TeBaQA, a QA system which learns to map question to SPARQL template mappings using basic graph pattern isomorphisms. TeBaQA significantly eases the domain/KB adoption process as it relies only on a benchmark dataset at its core.

In the future, we will evaluate TeBaQA on more heterogeneous KG benchmarks to identify further choke points. We will also improve the question classification and information extraction by using novel deep learning mechanism.

**Acknowledgements** We acknowledge the support of the Federal Ministry for Economic Affairs and Energy (BMWi) project SPEAKER (FKZ 01MK20011A), ScaDS.AI (01/S18026A) as well as the Fraunhofer Zukunftsstiftung project JOSEPH and the Eurostars Project PORQUE (FKZ 01QE2056C). This work was partially supported by the German Federal Ministry of Transport and Digital Infrastructure (BMVI) in the project LIMBO (no. 19F2029I) and by the German Federal Ministry of Education and Research (BMBF) in the project SOLIDE (no. 13N14456) within 'KMU-innovativ: Forschung für die zivile Sicherheit' in particular 'Forschung für die zivile Sicherheit'.

## References

- [1] A. Abujabal, R. S. Roy, M. Yahya, and G. Weikum. Never-ending learning for open-domain question answering over knowledge bases. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW*, pages 1053–1062, 2018.
- [2] A. Abujabal, M. Yahya, M. Riedewald, and G. Weikum. Automated template generation for question answering over knowledge graphs. In *Proceedings of the 26th International Conference on World Wide Web, WWW '17*, pages 1191–1200, 2017.
- [3] R. G. Athreya, S. Bansal, A. N. Ngomo, and R. Usbeck. Template-based question answering using recursive neural networks. *CoRR*, abs/2004.13843, 2020.
- [4] C. V. S. Avila, W. Franco, J. G. R. Maia, and V. M. P. Vidal. CONQUEST: A framework for building template-based IQA chatbots for enterprise knowledge graphs. In *NLDB*, volume 12089, pages 60–72. Springer, 2020.

- [5] J. Berant, A. Chou, R. Frostig, and P. Liang. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, 2013.
- [6] N. Chakraborty, D. Lukovnikov, G. Maheshwari, P. Trivedi, J. Lehmann, and A. Fischer. Introduction to neural network based approaches for question answering over knowledge graphs. *CoRR*, abs/1907.09361, 2019.
- [7] W. Cui, Y. Xiao, H. Wang, Y. Song, S. Hwang, and W. Wang. KBQA: learning question answering over QA corpora and knowledge bases. *PVLDB*, 10(5):565–576, 2017.
- [8] J. Daiber, M. Jakob, C. Hokamp, and P. N. Mendes. Improving efficiency and accuracy in multilingual entity extraction. In *Proceedings of the 9th International Conference on Semantic Systems (I-Semantics)*, 2013.
- [9] D. Diefenbach, A. Both, K. Singh, and P. Maret. Towards a question answering system over the semantic web. *Semantic Web*, 11(3):421–439, 2020.
- [10] D. Diefenbach, V. López, K. D. Singh, and P. Maret. Core techniques of question answering systems over knowledge bases: a survey. *Knowl. Inf. Syst.*, 55(3):529–569, 2018.
- [11] M. Dubey, D. Banerjee, A. Abdelkawi, and J. Lehmann. Lc-quad 2.0: A large dataset for complex question answering over wikidata and dbpedia. In *Proceedings of the 18th International Semantic Web Conference (ISWC)*. Springer, 2019.
- [12] F. Eibe, M. Hall, I. Witten, and J. Pal. The weka workbench. *Online Appendix for “Data Mining: Practical Machine Learning Tools and Techniques”*, 4, 2016.
- [13] O. Gervasi and V. Kumar. *Computational Science and Its Applications - ICCSA 2006: International Conference, Glasgow, UK, May 8-11, 2006, Proceedings*. Computational Science and Its Applications: ICCSA 2006. Springer, 2006.
- [14] Y. Gu, S. Kase, M. Vanni, B. M. Sadler, P. Liang, X. Yan, and Y. Su. Beyond I.I.D.: three levels of generalization for question answering on knowledge bases. *CoRR*, abs/2011.07743, 2020.
- [15] Y. Hao, H. Liu, S. He, K. Liu, and J. Zhao. Pattern-revising enhanced simple question answering over knowledge bases. In *Proceedings of the 27th International Conference on Computational Linguistics, COLING*, pages 3272–3282, 2018.
- [16] K. Höffner, S. Walter, E. Marx, R. Usbeck, J. Lehmann, and A. N. Ngomo. Survey on challenges of question answering in the semantic web. *Semantic Web*, 8(6):895–920, 2017.
- [17] P. Kapanipathi, I. Abdelaziz, S. Ravishankar, S. Roukos, A. G. Gray, R. F. Astudillo, M. Chang, C. Cornelio, S. Dana, A. Fokoue, D. Garg, A. Gliozzo, S. Gurajada, H. Karanam, N. Khan, D. Khandelwal, Y. Lee, Y. Li, F. P. S. Luus, N. Makondo, N. Mihindukulasooriya, T. Naseem, S. Neelam, L. Popa, R. G. Reddy, R. Riegel, G. Rossiello, U. Sharma, G. P. S. Bhargav, and M. Yu. Question answering over knowledge bases by leveraging semantic parsing and neuro-symbolic reasoning. *CoRR*, abs/2012.01707, 2020.
- [18] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, and C. Bizer. Dbpedia - A large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195, 2015.
- [19] D. Moussallem, R. Usbeck, M. Röder, and A. N. Ngomo. Entity linking in 40 languages using MAG. In *ESWC*, pages 176–181, 2018.
- [20] D. Moussallem, R. Usbeck, M. Röder, and A.-C. N. Ngomo. Mag: A multilingual, knowledge-base agnostic and deterministic entity linking approach. In *Proceedings of the Knowledge Capture Conference*, pages 1–8, 2017.
- [21] J. Pérez, M. Arenas, and C. Gutierrez. Semantics and complexity of sparql. *ACM Trans. Database Syst.*, 34(3):16:1–16:45, Sept. 2009.
- [22] M. Saleem, S. N. Dastjerdi, R. Usbeck, and A. N. Ngomo. Question answering over linked data: What is difficult to answer? what affects the F scores? In *2nd International Workshop on Benchmarking Linked Data and NLIWoD3*. CEUR-WS.org, 2017.
- [23] K. Singh, A. Both, A. S. Radhakrishna, and S. Shekarpour. Frankenstein: A platform enabling reuse of question answering components. In *ESWC*, volume 10843 of *Lecture Notes in Computer Science*, pages 624–638. Springer, 2018.
- [24] T. Soru, E. Marx, A. Valdestilhas, D. Esteves, D. Moussallem, and G. Publio. Neural machine translation for query construction and composition. *CoRR*, abs/1806.10478, 2018.
- [25] P. Trivedi, G. Maheshwari, M. Dubey, and J. Lehmann. Lc-quad: A corpus for complex question answering over knowledge graphs. In *International Semantic Web Conference*, pages 210–218. Springer,

- 2017.
- [26] C. Unger, L. Bühmann, J. Lehmann, A. N. Ngomo, D. Gerber, and P. Cimiano. Template-based question answering over RDF data. In *Proceedings of the 21st World Wide Web Conference 2012, WWW*, pages 639–648, 2012.
  - [27] R. Usbeck, R. H. Gusmita, A. N. Ngomo, and M. Saleem. 9th challenge on question answering over linked data (QALD-9). In *Joint proceedings of (SemDeep-4) and (NLIWOD-4)*, volume 2241 of *CEUR Workshop Proceedings*, pages 58–64, 2018.
  - [28] R. Usbeck, A. N. Ngomo, F. Conrads, M. Röder, and G. Napolitano. 8th challenge on question answering over linked data (QALD-8). In *Joint proceedings of (SemDeep-4) and (NLIWOD-4)*, volume 2241 of *CEUR Workshop Proceedings*, pages 51–57, 2018.
  - [29] R. Usbeck, A. N. Ngomo, B. Haarmann, A. Krithara, M. Röder, and G. Napolitano. 7th open challenge on question answering over linked data (QALD-7). In *ESWC*, pages 59–69, 2017.
  - [30] R. Usbeck, M. Röder, M. Hoffmann, F. Conrads, J. Huthmann, A. N. Ngomo, C. Demmler, and C. Unger. Benchmarking Question Answering Systems. *Semantic Web*, 10(2):293–304, 2019.
  - [31] S. Vakulenko, J. D. F. Garcia, A. Polleres, M. de Rijke, and M. Cochez. Message passing for complex question answering over knowledge graphs. In W. Zhu, D. Tao, X. Cheng, P. Cui, E. A. Rundensteiner, D. Carmel, Q. He, and J. X. Yu, editors, *CIKM*, pages 1431–1440. ACM, 2019.
  - [32] Y. Wang, J. Berant, and P. Liang. Building a semantic parser overnight. In *ACL*, pages 1332–1342, 2015.
  - [33] X. Yin, D. Gromann, and S. Rudolph. Neural machine translating from natural language to SPARQL. *CoRR*, abs/1906.09302, 2019.
  - [34] W. Zheng and M. Zhang. Question answering over knowledge graphs via structural query patterns. *CoRR*, abs/1910.09760, 2019.
  - [35] L. Zou, R. Huang, H. Wang, J. X. Yu, W. He, and D. Zhao. Natural language question answering over RDF: a graph data driven approach. In *International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014*, pages 313–324, 2014.