

# MQALD: Evaluating the impact of modifiers in Question Answering over Knowledge Graphs

Lucia Siciliani <sup>a,\*</sup>, Pierpaolo Basile <sup>a</sup>, Pasquale Lops <sup>a</sup>, and Giovanni Semeraro <sup>a</sup>

<sup>a</sup> *Department of Computer Science, University of Bari Aldo Moro, Via E. Orabona 4, 70125, Bari, Italy*  
E-mails: [lucia.siciliani@uniba.it](mailto:lucia.siciliani@uniba.it), [pierpaolo.basile@uniba.it](mailto:pierpaolo.basile@uniba.it), [pasquale.lops@uniba.it](mailto:pasquale.lops@uniba.it), [giovanni.semeraro@uniba.it](mailto:giovanni.semeraro@uniba.it)

**Editor:** Harald Sack, FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Germany

**Solicited reviews:** Simon Gottschalk, L3S Research Center, Leibniz Universität Hannover, Germany; Ricardo Usbeck, Data Science Group, Paderborn University / Fraunhofer IAIS, Germany; Rony Hasan, Fraunhofer IAIS, Germany; Two Anonymous reviewers

**Abstract.** Question Answering (QA) over Knowledge Graphs (KG) aims to develop a system that is capable of answering users' questions using the information coming from one or multiple Knowledge Graphs, like DBpedia, Wikidata, and so on. Question Answering systems need to translate the user's question, written using natural language, into a query formulated through a specific data query language that is compliant with the underlying KG. This translation process is already non-trivial when trying to answer simple questions that involve a single triple pattern. It becomes even more troublesome when trying to cope with questions that require modifiers in the final query, i.e., aggregate functions, query forms, and so on. The attention over this last aspect is growing but has never been thoroughly addressed by the existing literature. Starting from the latest advances in this field, we want to further step in this direction. This work aims to provide a publicly available dataset designed for evaluating the performance of a QA system in translating articulated questions into a specific data query language. This dataset has also been used to evaluate three QA systems available at the state of the art.

**Keywords:** Question Answering, Knowledge Graphs, Dataset, Benchmark, SPARQL modifiers

## 1. Introduction

Question Answering (QA) is a challenging task that allows users to acquire information from a data source through questions written using natural language. The user remains unaware of the underlying structure of the data source, which could be anything ranging from a collection of textual documents to a database. For this reason, QA systems have always represented a goal for researchers as well as industries since it allows the creation of a Natural Language Interface (NLI) [1, 2]. NLIs can make accessible, especially to non-expert

users, a huge amount of data that would have been undisclosed otherwise.

Several criteria can be used to categorize QA systems like the application domain (closed or open) or the kind of question that the system can handle (factual, causal, and so on) [3]. Nevertheless, one of the most distinguishing aspects is the structure of the data source used to retrieve the answers. This feature deeply affects the QA system structure and the methods applied to retrieve the correct answer.

The first QA systems were built as an attempt to create NLIs for databases [4] thus QA over structured data has been investigated since the late sixties. However, this research area remains a challenge, and the recent advances in the Semantic Web have brought more at-

---

\*Corresponding author. E-mail: [lucia.siciliani@uniba.it](mailto:lucia.siciliani@uniba.it).

1 attention to this topic. QA system can exploit the infor-  
 2 mation contained within Knowledge Graphs (e.g., DB-  
 3 pedia [5], or Wikidata [6]), to answer questions regard-  
 4 ing several topics and also can ease the access to this  
 5 great amount of data to non-expert users.

6 RDF<sup>1</sup> and SPARQL<sup>2</sup> represent the two standards  
 7 chosen by the W3C for data description language and  
 8 data query language for information resources of the  
 9 Web. Expressing an information need using SPARQL  
 10 to retrieve specific information is undoubtedly a non-  
 11 trivial task, especially for users that do not know this  
 12 kind of language. For this reason, QA systems have  
 13 to fulfill the role of interpreters, which take charge  
 14 of the translation of the users' information needs into  
 15 SPARQL queries.

16 However, translating the natural language, which is  
 17 inherently ambiguous, into a formal query using a data  
 18 query language like SPARQL is not easy. The main  
 19 problem that makes this process particularly difficult  
 20 is recognized in the literature as *lexical gap*. The lexical  
 21 gap indicates the syntactic distance between the natu-  
 22 ral language question and the correspondent SPARQL  
 23 query. In order to bridge it, it is necessary to perform  
 24 several Natural Language Processing tasks.

25 The first one is Entity Linking, i.e., find a match be-  
 26 tween portions of the question to entities within the  
 27 KG. This process can be straightforward if there is  
 28 a string matching between the question and the label  
 29 of a KG resource: in the question “*Which films were  
 30 produced by Steven Spielberg?*”, *Steven Spielberg*  
 31 can easily be mapped to its resource, e.g. `wiki:Q8877`  
 32 in Wikidata or `dbr:Steven_Spielberg` in DB-  
 33 pedia through the label of the resources. However,  
 34 many factors can easily make this task nontrivial. For  
 35 example, there could be a periphrasis like in the ques-  
 36 tion: “*When was the President of the USA born?*”.

37 Another important issue is related to the ambiguity  
 38 of the natural language. In the question “*Who directed  
 39 Philadelphia?*”, the word *Philadelphia* could refer to  
 40 both the film or the city. In such cases, it is necessary to  
 41 exploit the context provided by the question to select  
 42 the right resource.

43 Similar issues can be found when dealing with Re-  
 44 lation Linking, which requires a map between seg-  
 45 ments of the question and relations in the KG. Here  
 46 the lexical gap can be even wider since there can  
 47

48 <sup>1</sup><https://www.w3.org/TR/rdf11-concepts/>

49 <sup>2</sup><https://www.w3.org/TR/rdf-sparql-query/>

50 <sup>3</sup><https://www.wikidata.org/wiki/>

51 <sup>4</sup><http://dbpedia.org/resource/>

1 be many expressions that refer to a particular re-  
 2 lation. Questions like “*Who is the wife of Barack  
 3 Obama?*”, “*Who is the spouse of Barack Obama?*”,  
 4 “*Who married Barack Obama?*” always refer to the  
 5 same relation, which is `dbo:spouse` in DBpedia  
 6 and `wiki:Property:P26` in Wikidata.

7 These two issues can be combined in several ways.  
 8 An interesting example is represented by requests like  
 9 “*Give me all Dutch parties*”. Here, to find the answer,  
 10 the phrase *Dutch parties* requires a combination of En-  
 11 tity Linking and Relation Extraction, and the relation  
 12 is implicit. Moreover, not one but two relations need to  
 13 be concatenated to obtain the final query since we are  
 14 looking for something that is a political party (first re-  
 15 lation), and it has to be based in the Netherlands (sec-  
 16 ond relation).

17 These two problems, intertwined, represent a strug-  
 18 gle for any QA system, and several techniques have  
 19 been applied in the literature to overcome them. How-  
 20 ever, there is another kind of lexical gap that has to be  
 21 considered. For example, to find an answer for ques-  
 22 tions like “*Which presidents were born in 1945?*”,  
 23 mapping phrases of the question to the right resources  
 24 in the KG is not sufficient: the question must be trans-  
 25 lated into a query that involves the proper SPARQL  
 26 modifier. By SPARQL modifiers, we refer to all the  
 27 constructs that alter the basic SPARQL structure made  
 28 up of a SELECT clause followed by a WHERE clause  
 29 that encloses a graph pattern composed of one or more  
 30 triple patterns. In the previous example, the question  
 31 “*Which presidents were born in 1945?*” has to be  
 32 translated into a query containing the FILTER modi-  
 33 fier like shown in Listings 1.

```
34
35 SELECT DISTINCT ?uri WHERE {
36   ?uri a dbo:President.
37   ?uri dbo:birthDate ?date.
38   FILTER regex(?date, '^1945') }
39
```

40 Listing 1: SPARQL query for “*Which presidents were  
 41 born in 1945?*”

42  
 43 The SPARQL query language makes available sev-  
 44 eral modifiers: query forms (like ASK), solution se-  
 45 quence modifiers (e.g. ORDER BY, LIMIT, OFFSET),  
 46 functions (e.g. FILTER, COUNT, SUM, AVG, NOW,  
 47 YEAR, MONTH). These modifiers can be used alone for  
 48 simpler queries, but more frequently, they have to be  
 49

50 <sup>5</sup><http://dbpedia.org/ontology/>

1 combined to obtain a SPARQL query that is correct  
2 and allows to retrieve all the right answers.

3 The construction of queries containing modifiers  
4 still represents an issue that has been explicitly ex-  
5 plored in literature only by few works [7–12], due to  
6 the complexity of the task and the lack of resources  
7 that can help researchers to develop and evaluate novel  
8 solutions. With this paper, we want to step forward in  
9 this direction and propose MQALD: a dataset created  
10 to allow KGQA systems to take on this challenge and  
11 overcome the issues related to this specific research  
12 area.

13 The paper is organized as follows: in Section 2  
14 we describe the state-of-the-art datasets and evalua-  
15 tion metrics used in this research area; in Section 3  
16 we introduce our MQALD dataset, explaining how it  
17 was created and examining the modifiers it contains; in  
18 Section 4 we introduce the QA systems chosen for the  
19 evaluation and describe how it was performed; in Sec-  
20 tion 5 we analyze the results obtained by the chosen  
21 systems, and finally, in Section 6, we draw the conclu-  
22 sion and propose some insights for future research.

## 24 2. Related Work

25  
26  
27 The development of a unified dataset for evaluating  
28 KGQA systems started in 2011 with the first edition of  
29 the Question Answering over Linked Data (QALD)<sup>6</sup>  
30 challenge. Within each edition of the challenge, the or-  
31 ganizers always proposed different tasks, the most im-  
32 portant being QA over DBpedia. The first dataset pro-  
33 posed for this task was composed of 100 questions,  
34 equally split into a train and test set. After each edi-  
35 tion, this set of questions was modified and further ex-  
36 panded, leading to 408 questions in the training set  
37 and 150 questions in the test set for the last edition of  
38 the challenge. QALD datasets contain questions that  
39 are manually compiled and curated to include different  
40 grades of complexity. Every dataset is composed of a  
41 set of English questions that (since QALD-3) is also  
42 translated into several languages, like German, Span-  
43 ish, Italian, and French. Each question is associated  
44 with its answer (which is a set composed of one or  
45 more resources or literal values) and, most importantly,  
46 also the SPARQL query that translates the input ques-  
47 tion.

48 Besides the QALD dataset, other datasets have been  
49 proposed to benchmark QA systems over KGs.

50  
51 <sup>6</sup><http://qald.aksw.org/>

1 WebQuestions and SimpleQuestion were both built  
2 for Freebase [13]. WebQuestions [14] contains 5,810  
3 questions obtained by exploiting the Google Suggest  
4 API and the Amazon Mechanical Turk. Due to this  
5 generation process, questions within this dataset usu-  
6 ally follow similar templates, and each of them re-  
7 volves around a single entity of the KG. All questions  
8 are associated with an entity of Freebase that can be  
9 found in the question and the right answer.

10 SimpleQuestions [15] instead is the dataset with the  
11 highest number of questions, and it was created to pro-  
12 vide a richer dataset. It consists of 108,442 questions  
13 that have been created in two steps: in the first one,  
14 a list of facts or triples were extracted directly from  
15 Freebase, then these triples were sent to human annota-  
16 tors who were in charge of translating such triples into  
17 natural language questions. Like in the QALD dataset,  
18 each question is associated with its SPARQL version.  
19 The dataset derives its name from the fact that all the  
20 questions can be answered by exploiting just one rela-  
21 tion of the KG.

22 LC-QuAD [16] is based on Wikidata using tem-  
23 plates that allow formulating questions that are trans-  
24 lated into SPARQL queries that require 2-hop rela-  
25 tions. These templates are then reviewed by human an-  
26 notators that verify the correctness of the questions.  
27 The final dataset comprises 5,000 questions, with the  
28 corresponding SPARQL translations and the templates  
29 used for generating them. With LC-QuAD 2.0 [17] the  
30 dataset has been further extended and now contains  
31 30,000 questions.

32 TempQuestion [18] is a dataset based on Free-  
33 base and obtained by collecting questions contain-  
34 ing temporal expressions from already existing source  
35 datasets. ComplexWebQuestions [19] is also based on  
36 Freebase and requires reasoning over multiple web  
37 snippets. Finally Event-QA [20] is a novel dataset  
38 based on a custom KG called EventKG<sup>7</sup>.

39 As aforementioned, QALD represents a dataset and  
40 a challenge that makes available a benchmark for QA  
41 systems. Factoid QA inherited its evaluation metrics  
42 from classical evaluation systems, thus estimating the  
43 performance in terms of Precision, Recall, and F-  
44 measure. Since the seventh edition of the challenge,  
45 the GERBIL benchmark platform [21, 22] was used to  
46 evaluate the performance of the participating systems.  
47 GERBIL, which is publicly available online<sup>8</sup>, embeds  
48 several QA systems and allows to evaluate different

50 <sup>7</sup><http://eventkg.l3s.uni-hannover.de/>

51 <sup>8</sup><http://gerbil-qa.aksw.org/gerbil/>

QA datasets, including those distributed during all the QALD challenges.

### 3. MQALD: Dataset Creation

In this section, we introduce MQALD, a dataset that is composed only of questions containing modifiers. Two sections compose MQALD: the first one contains novel questions, and the second one is obtained by extracting questions that require modifiers from the QALD dataset. In the following, we will describe with more detail each portion of MQALD.

#### 3.1. Generating novel questions

MQALD comprises 100 questions that have been manually created by human annotators based on DBpedia 2016-10. To create MQALD, we employed two annotators who are both familiar with the use of SPARQL and SQL-like Data Query Languages. Since the dataset aims to evaluate QA systems' capability to translate complex information needs into query with modifiers, we did not apply any constraint to the question's creation process. Annotators are free to choose question topics and modifiers. In this way, the results obtained by a system are more likely to depend on its ability to create questions with modifiers.

Modifier	#occurrences
LIMIT	25
ORDER BY	31
FILTER	41
ASK	20
UNION	10
OFFSET	6
COUNT	26
GROUP BY	15
HAVING	10
YEAR	3
NOW	0

Table 1

Frequencies of each modifier within the novel questions available in MQALD.

Each annotator was in charge of developing 50 questions and then review the other 50 questions created by the other annotator.

Each question is provided in four languages: English, Italian, French, and Spanish. Both annotators have Italian as their mother tongue and have a good

knowledge of the English language. Each question was first written in English and then translated to Italian. For the other two languages, translations of the original English question were obtained through machine translation<sup>9</sup>.

```
{
  "id": 187,
  "aggregation": true,
  "hybrid": false,
  "question": [
    {
      "string": "Which is the fifth most
        populous country in the world?",
      "keywords": "fifth, most populous,
        country",
      "language": "en",
      ...
    }
  ],
  "answertype": "resource",
  "onlydbo": true,
  "query": {
    "sparql": "PREFIX dbo: <http://
      dbpedia.org/ontology/> PREFIX
      rdf: <http://www.w3.org/1999/02/
      22-rdf-syntax-ns#> SELECT ?
      country WHERE { ?country rdf:
      type dbo:Country; dbo:
      populationTotal ?population }
      ORDER BY DESC(xsd:integer(?
      population)) LIMIT 1 OFFSET 4"
    },
    "answers": [
      { "head": { "vars": ["country"] } },
      "results": {
        "bindings": [
          { "country": { "type": "uri", "value":
            "http://dbpedia.org/resource/
            China" } }
        ]
      }
    ],
    "modifiers": ["ORDER BY", "LIMIT", "
      OFFSET"]
  }
}
```

Listing 2: JSON structure of a question within the dataset.

The two annotators discussed any further discrepancy until they met an agreement<sup>10</sup>. Questions were

<sup>9</sup>We use Google Translate.

<sup>10</sup>During the annotation process, the first annotator edited the 12% of questions of the second annotator, while the 18% of questions of the first annotator were rectified.

built paying attention so that the modifiers occurring in each SPARQL query were:

- mandatory: the use of the modifier is indispensable to retrieve the right answer;
- KG independent: modifiers are not used to adapt the query to a particular structure of the underlying KG. An example of KG dependency is represented by the use of the `LIMIT` modifier to retrieve a specific resource or literal from the result set. An example of this behavior is given in Section 3.3.3.

Considering these two aspects is particularly important since, in this way, it is possible to guarantee that each modifier is connected to a specific semantics of the natural language question. Table 1 shows the distribution of modifiers among the 100 questions made by the annotators.

The MQALD dataset is in JSON format and compliant with the QALD dataset structure: each question is serialized as a JSON object which is then stored in a JSON array that contains all the questions as shown in Listings 2. The only exceptions in the structure of each question is represented by the presence of a further JSON array, named `modifiers`, containing all the modifiers occurring in the SPARQL query (this array is empty if there are no modifiers).

### 3.2. Extracting queries with modifiers from QALD

To further extend our dataset, we also collect questions with modifiers available in the QALD dataset. As stated in Section 2, in each edition of QALD, the dataset is obtained as an extension of the previous edition dataset, although we noticed there was not a complete overlap. For this reason, we decided to merge the questions contained in the last three editions of the QALD challenge over DBpedia: QALD-7 [23], QALD-8 [24], and QALD-9 [25]. We decide not to include queries with modifiers from other datasets since they are not based on DBpedia or are based on a non-standard version of DBpedia<sup>11</sup>.

Considering that each QALD dataset is split in training and test, the merge of the three editions was done separately. Therefore two lists of questions were obtained: one containing all the training questions from QALD-9, 8, and 7, and the other one containing all

<sup>11</sup>This is the case of LC-QuAD that uses a version of DBpedia, called DBpedia 2018, that merges DBpedia with information coming from Wikidata.

the test questions from QALD-9, 8, and 7. Since training data could have been used to develop the systems, we excluded from the test set those questions that also appeared within the training data in any of the three editions. Thus, those questions were removed from the test questions of QALD-9, 8, and 7 and added to the list composed by the training set questions. We obtained the following files:

- **QALD\_train**: contains the merge of all questions coming from the training sets of QALD-7, QALD-8, QALD-9;
- **QALD\_train\_MODS**: contains only the questions that require a modifier extracted from QALD\_train;
- **QALD\_train\_NO\_MODS**: contains only the questions that do not require a modifier extracted from QALD\_train;
- **QALD\_test**: contains the merge of all questions coming from the test sets of QALD-7, QALD-8, QALD-9;
- **QALD\_test\_MODS**: contains only the questions that require a modifier extracted from QALD\_test;
- **QALD\_test\_NO\_MODS**: contains only the questions that do not require a modifier extracted from QALD\_test;

Dataset	#questions
QALD_train	462
QALD_train_MODS	154
QALD_train_NO_MODS	308
QALD_test	115
QALD_test_MODS	41
QALD_test_NO_MODS	74

Table 2

Number of questions included in each file.

Table 2 shows the number of questions included in each file. The format of the JSON file remains the same as the one shown in Listings 2 with the addition of the attribute `qald-version` which indicates the QALD dataset of origin. Table 3 shows the distribution of modifiers within the training and test sets obtained by merging QALD-9, 8, and 7.

### 3.3. Modifiers within the dataset

In this section, we will describe in detail the modifiers that are contained in the dataset through real examples extracted from MQALD. The modifiers will be listed according to their frequency.

Modifier	QALD train	QALD test
LIMIT	48	11
ORDER BY	45	9
FILTER	34	16
ASK	40	3
UNION	33	9
OFFSET	33	5
COUNT	29	7
GROUP BY	3	2
HAVING	3	1
YEAR	2	2
NOW	0	2

Table 3

Frequencies of each modifier obtained after merging the last three editions of the QALD dataset.

### 3.3.1. FILTER

The `FILTER` keyword is used to restrict the number of results that match the graph pattern specified just before the `FILTER` itself. The `FILTER` keyword precedes an expression that can be constituted by a wide variety of operators categorized according to the number of parameters involved.

There are *unary operators* like “!” which represents the logical NOT and “BOUND” that checks if its argument is bounded to a value. These operators can be used alone or in conjunction to answer questions like “*Is Frank Herbert still alive?*”. In this case, the answer can be found only checking if the object of the relation `res:Frank_Herbert dbo:deathDate ?date` does not exist in the KG (i.e. `FILTER(!BOUND(?date))`).

We also have *binary operators* that comprise logical connectives AND “&&” and OR “||”, as well as all the other test operators between two values (like `==`, `!=`, `>`, `<`, and so on). These are used to answer questions like “*Which caves have more than 3 entrances?*” creating a SPARQL query as shown by Listing 3.

```
SELECT DISTINCT ?uri WHERE {
  ?uri a dbo:Cave;
  dbp:entranceCount ?entrance
  FILTER (?entrance>3)}
```

Listing 3: SPARQL query for “*Which caves have more than 3 entrances?*”.

Finally, there is also a single *ternary operator* which is `REGEX`, that checks if a certain string (first param-

ter) matches a regular expression (second parameter). The third parameter is an optional flag, which specifies a case-insensitive pattern match). This ternary operator can be used in questions like “*Is the wife of president Obama called Michelle?*” to check if the name “*Michelle*” is contained within the label of the resource `dbo:Michelle_Obama`.

### 3.3.2. ORDER BY

The second most common modifier is `ORDER BY`, which is used to alter the order of the solution sequence. It is important to notice that the dataset, the `ORDER BY` modifier is always used in conjunction with other modifiers except one question within the dataset. An example where `ORDER BY` appears by itself is “*What are the top selling luxury vehicle brands in Germany?*”, where the answer is represented by all the results of an ordered list.

### 3.3.3. LIMIT

The `LIMIT` modifier is usually used in combination with `ORDER BY` and `OFFSET` which allows answering questions containing superlatives, i.e. “*Who is the tallest player of the Atlanta Falcons?*” or “*For which label did Elvis record his first album?*”.

Another combination is composed of the previous one (`LIMIT`, `ORDER BY`, and `OFFSET`) with the addition of the `COUNT` aggregation function for questions like “*Which country has the most official languages?*” where the `ORDER BY` needs to be applied to the object of the `COUNT` function as illustrated by Listing 4.

```
SELECT DISTINCT ?uri WHERE {
  ?uri a dbo:Country;
  dbp:officialLanguages ?language.}
ORDER BY DESC(COUNT(?language))
OFFSET 0 LIMIT 1
```

Listing 4: SPARQL query for “*Which country has the most official languages?*”.

The `LIMIT` modifier is also subject of controversial use within the QALD dataset. An example is represented by the question “*With how many countries Iran has borders?*” that translates into a SPARQL query with the `LIMIT` modifier set to 8. This is correct since by executing the query against DBpedia without this modifier, only the first 8 results represent the right answer (i.e. the resources corresponding to the countries Iran has borders with), while the following are just literals not relevant to the question. However, in this case,

the usage of the `LIMIT` modifier is deeply connected to the structure of DBpedia.

For the novel questions introduced in MQALD we employed the `LIMIT` modifier to translate the semantics of superlatives or select elements in a specific position from an ordered list as shown in Listing 5.

```
SELECT ?capital WHERE {
  ?uri dbo:capital ?capital;
  dbo:populationDensity ?density}
ORDER BY DESC(?density)
LIMIT 2 OFFSET 2
```

Listing 5: SPARQL query for “Which are the second and third capitals in the world with higher population density?”.

### 3.3.4. ASK

Among all the possible SPARQL query forms besides `SELECT`, the only one present within the dataset is `ASK` which is reasonable since QA systems on KGs are usually factoid QA systems. The `ASK` query form is necessary to construct questions that require a Yes/No answer thus returning a Boolean value, like “Did Socrates influence Aristotle?”.

```
ASK WHERE {
  res:Aristotle
  dbo:influencedBy res:Socrates }
```

Listing 6: SPARQL query for “Did Socrates influence Aristotle?”.

### 3.3.5. COUNT

Besides from being used in combination with `ORDER BY` and `LIMIT`, the `COUNT` modifier is used alone, for example in questions that start with “How many”, like “How many films did Leonardo DiCaprio star in?”.

### 3.3.6. UNION

The `UNION` modifier is used to merge the results of two different patterns. Within the QALD dataset, it is usually employed to take into account more relations for obtaining the full set of answers as shown in Listing 7.

Despite being correct, the use of the `UNION` modifier in this query is associated with how the KG is structured, i.e., it is necessary to know that for retrieving all the companies that are in Munich, we need three

relations (`dbo:location`, `dbo:headquarter`, and `dbo:locationCity`).

For this reason, in the questions we introduced in MQALD, through the help of human annotators, we employed the `UNION` modifier as a translator of the semantics of the logic `OR` like shown in Listings 8.

```
SELECT DISTINCT ?uri WHERE {
  ?uri a dbo:Company
  {?uri dbo:location dbr:Munich}
  UNION
  {?uri dbo:headquarter dbr:Munich}
  UNION
  {?uri dbo:locationCity dbr:Munich}}
```

Listing 7: SPARQL query for “Give me all companies in Munich.”.

```
SELECT ?uri WHERE {
  ?uri rdf:type dbo:Band.
  ?uri dbo:genre dbr:Electronic_music.
  {{?uri dbo:hometown dbr:France}
  UNION
  {?uri dbo:hometown dbr:Italy}}
```

Listing 8: SPARQL query for “Give me all the electronic music bands having hometown in France or in Italy.”.

### 3.3.7. OFFSET

The `OFFSET` modifier is used to shift the start of the retrieved solutions after the specified number.

In the questions extracted from QALD, the majority of times, the queries include an `OFFSET` of 0, which has no effect on the results and therefore could also be neglected. However, there are some questions where there is an `OFFSET` of a different number that allows answering correctly to the input question but appears controversial. Let us consider the question “How many foreigners speak German?”. This question is translated into a SPARQL query where `OFFSET 1` is applied over the object of the relation `dbr:German_language dbp:speakers ?Ger_lang`. This relation has as object three literal values, namely: “L2 speakers: 10–15 million”, “as a foreign language: 75–100 million”, and “million to 95 million”, so actually the use of `OFFSET 1` allows to pick up the second literal which represents the right answer. Nevertheless, obtaining this

SPARQL query without knowing how the KG is structured is infeasible, especially considering that in this case (but this frequently happens in DBpedia), the property `dbp:speakers` does not have any specification concerning its domain and range thus, the object of this property is not consistent (there can be strings, numbers or even resources). A similar scenario occurs for the question “*Where is the most deep point in the ocean?*” where a `OFFSET 13` is required to retrieve the answer.

The other existing case of non-zero `OFFSET` within the QALD questions ensues from the question “*What is the second highest mountain on Earth?*”, where the semantics of the question justifies the use of `OFFSET 1`.

In our dataset, we always employ the `OFFSET` modifier with the aforementioned semantics. `OFFSET` is thus used in conjunction with `ORDER BY` and `LIMIT` to obtain a particular element of an ordered list.

```
SELECT DISTINCT ?uri WHERE {
  ?uri a dbo:Mountain;
  dbo:elevation ?elevation}
ORDER BY DESC(?elevation)
OFFSET 1 LIMIT 1
```

Listing 9: SPARQL query for “*What is the second highest mountain on Earth?*”.

### 3.3.8. GROUP BY - HAVING

The `GROUP BY` modifier is used to divide the result set into groups, which can then be used as input to an aggregation function like `COUNT`.

```
SELECT DISTINCT ?uri WHERE {
  ?uri a dbo:Country. ?cave a dbo:Cave
  { ?cave dbo:location ?uri } UNION
  { ?cave dbo:location ?loc .
  ?loc dbo:country ?uri } }
GROUP BY ?uri
HAVING ( COUNT(?cave) > 10 )
```

Listing 10: SPARQL query for “*Which countries have more than ten caves?*”.

If the query requires to filter the result of aggregation, then it must be introduced with the modifier `HAVING` (instead of the keyword `FILTER`), as it happens for the question “*Which countries have more than ten caves?*” in Listing 10.

### 3.3.9. YEAR - NOW

SPARQL makes available several functions for dates and times, even though `YEAR` and `NOW` are the only two occurring in the dataset. The `YEAR` function returns the year part of the date given as an argument, while the `NOW` function does not require an argument and returns the current time (formatted according to the XSD `dateTime` format). These functions can be used to answer very peculiar questions, like “*Give me all American presidents of the last 20 years.*”.

```
SELECT DISTINCT ?uri WHERE {
  ?uri rdf:type dbo:Person;
  dct:subject
  dbo:Presidents_of_the_United_States;
  dbo:activeYearsEndDate ?d
  FILTER ( (YEAR(NOW()) - YEAR(?d)) <= 20 ) }
```

Listing 11: SPARQL query for “*Give me all American presidents of the last 20 years.*”.

### 3.4. Dataset availability and usage statistics

The dataset can be downloaded from Zenodo<sup>12</sup> and was published under the MIT license. A description of the dataset is provided within the repository. The code used for creating the dataset and the evaluation is freely available on GitHub<sup>13</sup>. We published the first version of the dataset in May 2020. Since then, the dataset has collected 70 unique views and 26 unique downloads<sup>14</sup>.

## 4. Evaluation

We compare the performance of three state of the art QA systems for KGs over the different datasets we constructed. Our evaluation aims to examine how QA systems perform on questions that require one or more modifiers within the final SPARQL query and how this result impacts the performance of the system itself. The QA systems chosen for the evaluation are the following:

- **gAnswer**<sup>15</sup>: the system proposed by Hu et al. [26] that won the QALD-9 challenge. The method

<sup>12</sup><https://zenodo.org/record/4657496>

<sup>13</sup><https://github.com/lsciliani/MQALD>

<sup>14</sup>Statistics computed on 1st April 2021.

<sup>15</sup><http://ganswer.gstore-pku.com/api/qald.jsp>

used by this system consists of transforming the input question into a Semantic Query Graph, a particular graph where each node corresponds to a resource and each vertex to a relation of the underlying KG. Resources are found using DBpedia Lookup, while relations are extracted using a Multi-Channel Convolutional Neural Network. Next, the Semantic Query Graph is translated into a SPARQL question that is then used to obtain the final answer;

- **QAnswer**<sup>16</sup>: a QA system developed by Diefenbach et al. [27]. It represents the extension of the WDAqua-core0 system [28] that took part in the QALD-9 challenge gaining the second place. The approach behind QAnswer is based upon four steps: the first one consists of retrieving all the resources that can be linked to a question by considering its n-grams, the second phase uses these resources to create all the possible queries according to specific patterns in the third step these queries are then ranked according to several features and finally the query that obtained the highest score is executed, and the answer is returned to the user;
- **TeBaQA**<sup>17</sup>: created by Vollmers et al. [29], exploits a template-based approach. The SPARQL templates are generated upon the questions available for the QALD challenge and a classifier is trained to assign each question to a specific template. Given a question, the classifier assigns to it a template which is later filled to build the final SPARQL query. TeBaQA ranked third during the QALD-9.

We decided to implement an evaluation tool following the guidelines reported in the QALD-9 report [25]. In particular, the evaluation script computes for each query  $q$  a set of metrics, such as precision (P), recall (R), and F-measure according to the following equations:

$$P(q) = \frac{\#correct\ system\ answers\ for\ q}{\#system\ answers\ for\ q}$$

$$R(q) = \frac{\#correct\ system\ answers\ for\ q}{\#gold\ standard\ answers\ for\ q}$$

<sup>16</sup><http://qanswer-core1.univ-st-etienne.fr/api/gerbil>

<sup>17</sup><http://139.18.2.39:8187/>

$$F(q) = \frac{2 \times P(q) \times R(q)}{P(q) + R(q)}$$

Moreover, there are additional information to take into account:

- If the golden answer-set is empty and the system retrieves an empty answer, precision, recall and F-measure are equal to 1.
- If the golden answer-set is empty but the system responds with any answer-set, precision, recall, and F-measure are equal to 0.
- If there is a golden answer but the the QA system responds with an empty answer-set, it is assumed that the system could not respond. Then the precision, recall, and F-measure are equal to 0.
- In any other case, the standard precision, recall, and F-measure are computed.

In our paper, we consider the macro measures as the QALD challenge: we calculated precision, recall, and F-measure per question and averaged the values at the end. As adopted in QALD, for the final evaluation, the Macro F1 QALD metric is used. This metric uses the previously mentioned additional information with the following exception:

- If the golden answer-set is not empty but the QA system responds with an empty answer-set, it is assumed that the system determined that it cannot answer the question. Then the precision is set to 1 and the recall and F-measure to 0.

## 5. Results

The results of the evaluation of the three mentioned QA systems, i.e. gAnswer, QAnswer, and TeBaQA are reported in Table 4.

Overall, an analysis of the performance of the systems in terms of F1-QALD measure for the datasets without modifiers (no\_mods) and those containing all the questions (with and without modifiers) confirms the results reported by the QALD-9 challenge: gAnswer represents the best system at the state-of-the-art for this task, followed by QAnswer and TeBaQA.

However, this ranking changes when comparing the results containing only questions that require modifiers. In fact, in this case, QAnswer outperforms the results obtained by gAnswer. To obtain more insights regarding this aspect, we checked which questions that

	gAnswer				QAnswer				TeBaQA			
	P	R	F	F1-Q	P	R	F	F1-Q	P	R	F	F1-Q
qald-9_test	.607	.316	.296	<b>.416</b>	.459	.222	.197	.299	.644	.141	.139	.231
mqald	.497	.037	.031	.069	.550	.083	.056	<b>.145</b>	.531	.039	.046	.072
mqald (only new queries)	.472	.032	.032	.060	.320	.093	.071	<b>.144</b>	.483	.036	.045	.068
mqald+qald_test_NO_MODS	.526	.173	.162	<b>.260</b>	.394	.161	.137	.228	.562	.104	.109	.176
qald_train_all	.577	.334	.313	<b>.423</b>	.529	.334	.302	.409	.602	.154	.144	.246
qald_train_MODS	.350	.054	.027	.094	.356	.223	.175	<b>.274</b>	.479	.051	.047	.092
qald_train_NO_MODS	.474	.466	.456	<b>.470</b>	.612	.392	.369	.478	.657	.199	.186	.306
qald_test_all	.590	.295	.284	<b>.393</b>	.465	.232	.203	.309	.634	.155	.156	.249
qald_test_MODS	.608	.049	.054	.091	.391	.139	.087	<b>.205</b>	.671	.045	.048	.085
qald_test_NO_MODS	.580	.431	.412	<b>.495</b>	.515	.269	.260	.353	.606	.215	.215	.317

Table 4

Results of the three systems over the ten datasets. The metrics used are Precision (P), Recall (R), F-Measure (F), and F1-QALD measure (F1-Q).

require modifiers are answered by each system within the test set, and the results are shown in Table 5 for questions contained in the QALD testset and Table 6 for the novel questions added in MQALD. By analyzing Precision, Recall, and F-measure calculated for each question and the set of answers returned by each system, we can have an insight into how they handle modifiers.

Out of the 41 questions that require modifiers extracted from the QALD datasets, the number of questions where the F-score is not equal to zero is 3 for gAnswer, 7 for QAnswer, and 3 for TeBaQA. Most of the questions answered by the three systems involve the UNION modifier. As stated in Section 3.3, within the QALD datasets, the UNION modifier is exclusively used to merge the results of different triple patterns rather than encapsulating a particular semantics and sometimes are even unnecessary due to the constant evolution of DBpedia. This is the example of the question “Which software has been published by Mean Hamster Software?” where the UNION modifier is not needed at all since the query `SELECT DISTINCT ?uri WHERE {?uri onto:publisher res:Mean_Hamster_Software}` is sufficient to retrieve all the answers. This could be due to a mistake or a lack of update of the dataset and allows gAnswer and QAnswer to obtain the right answer without using the modifier. Another case is represented instead by questions like “Which countries are connected by the Rhine?” where the UNION modifier is well used since it is necessary to retrieve the whole set of answers provided by the QALD. In this case, gAnswer and TeBaQA manage to return an answer but only the precision is equal

to 1, meaning that not all the answers have been retrieved and UNION was not used.

The same happens for all the questions where the F-score is lower than 1. For example, QAnswer answers the question “What is the longest river in China?” by returning all the rivers that flow through China and, of course, the longest river is also included in this set. The use of the three modifiers (LIMIT, OFFSET, and ORDER BY) would have allowed to properly order the list by the rivers’ length and pick the first element of the ordered list. With this result, we can deduce that modifiers were not used to retrieve the answer.

Also, for the question “How many grand-children did Jacques Cousteau have?” we can observe a misleading behavior since the number of grand-children of Jacques Cousteau actually coincides with his children’s number. Thus QAnswer could answer correctly by just considering the property `dbp:children`. It is important to notice that the property `dbp:children` does not exist in the current version of DBpedia: we have the property `dbo:child` which has as object the resources of two of the four children of Jacques Cousteau. This implies that also this question should be updated properly to be compliant with the KG. Nevertheless, the property `dbp:children` is available in the 2016-10 version of DBpedia, and the object of this property is 4, which can be obtained without modifiers and erroneously used as the correct answer.

There is only one question that requires an ASK among all the others: “Is Pamela Anderson a vegan?”. Only QAnswer returns the right answer but using a SPARQL query that is completely unrelated to the question that just checks the existence of the resource `dbr:Pamela_Anderson`.

gAnswer					
id	Question	Modifiers	P	R	F
7	Which software has been published by Mean Hamster Software?	UNION	1	1	1
19	Give me all cars that are produced in Germany.	UNION	.920	.191	.316
24	Which countries are connected by the Rhine?	UNION	1	.833	.909
QAnswer					
id	Question	Modifiers	P	R	F
1	What is the highest mountain in Germany?	LIMIT, ORDER BY	.001	1	.002
5	Which airports are located in California, USA?	UNION	1	.351	.520
7	Which software has been published by Mean Hamster Software?	UNION	1	1	1
8	How many grand-children did Jacques Cousteau have?	COUNT	1	1	1
12	What is the longest river in China?	LIMIT, OFFSET, ORDER BY	.016	1	.032
15	Is Pamela Anderson a vegan?	ASK	1	1	1
31	Which politicians were married to a German?	UNION	.063	.091	.074
TeBaQA					
id	Question	Modifiers	P	R	F
6	Which countries in the European Union adopted the Euro?	UNION	.5	.022	.043
14	How many awards has Bertrand Russell?	COUNT	1	1	1
24	Which countries are connected by the Rhine?	UNION	1	.833	.909

Table 5

List of the questions answered by each system over the qald\_test\_MODS dataset.

An exception is represented by the question “*How many awards has Bertrand Russell?*” which requires the COUNT modifier since this information is not available in any other way from DBpedia.

Regarding the novel questions included in MQALD (new queries), 4 answers were provided by gAnswer, 11 by QAnswer, and 8 by TeBaQA. The only two systems that manage to return a full answer, with Precision, Recall, and F-measure at 1, are QAnswer and TeBaQA. As stated in Section 3, MQALD questions are formulated such that modifiers are necessary to retrieve the correct set of answers; thus, by analyzing these scores, it is possible to identify which modifiers are covered by which system easily. QAnswer is able to manage the ASK and COUNT modifiers while TeBaQA is capable to handle the COUNT modifiers as well. Of course, if a system can handle a specific modifier, it does not mean that it will answer all the questions containing that modifier since the lexical gap still represents a major problem in the translation from natural language to SPARQL.

Overall, the results show how all the systems under analysis do not perform well at handling modifiers, just with a few exceptions. Nevertheless, the capability of QAnswer to handle ASK and COUNT modifiers and better bridge the lexical gap by at least returning partial answers granted it the best score among the three selected systems.

## 6. Conclusions and Future Works

In this paper, we have introduced MQALD, a dataset of questions that contain SPARQL modifiers. The dataset is composed of 100 new manually created queries plus 41 queries with modifiers of the last three editions of the QALD. These modifiers were then analyzed to better understand their functioning considering the SPARQL syntax and how they can reflect specific semantics of the natural language query. Through the evaluation of three systems available at the state-of-the-art, emerged that there is still much work that must be done to create a QA system capable of han-

gAnswer					
id	Question	Modifiers	P	R	F
164	Give me all Stephen King books for which the publication date is known.	FILTER	.488	.976	.651
189	Give me the list of who directed and produced Saving Private Ryan.	UNION	.600	.750	.667
226	What are the first publication date and the last publication date of Mazinger Z?	UNION	1	.500	.667
243	Which Elvis songs are part of the Elvis is Back album?	FILTER, ORDER BY	.106	1	.192
QAnswer					
id	Question	Modifiers	P	R	F
162	What is the lowest mountain in Germany?	LIMIT, ORDER BY	.001	1	.002
169	Is Leonardo Da Vinci the author of the Mona Lisa?	ASK	1	1	1
170	What are the first 10 works in alphabetical order made by Leonardo da Vinci?	LIMIT, ORDER BY	.286	1	.444
182	Which films of the Jurassic Park saga has Steven Spielberg directed?	FILTER	.667	1	.800
183	How many movies has George Lucas directed?	COUNT	1	1	1
189	Give me the list of who directed and produced Saving Private Ryan.	UNION	1	.250	.400
194	Is Barack Obama a politician?	ASK	1	1	1
220	Which films in which Catherine Zeta-Jones starred in have earned less than the initial budget?	FILTER, ORDER BY	.077	1	.143
226	What are the first publication date and the last publication date of Mazinger Z?	UNION	1	.5	.667
231	Who are the author and publisher of Don Quixote?	UNION	1	.5	.667
238	How many cartoons are produced by Walt Disney?	COUNT	1	1	1
TeBaQA					
id	Question	Modifiers	P	R	F
170	What are the first 10 works in alphabetical order made by Leonardo da Vinci?	LIMIT, ORDER BY	.250	.1	.143
174	Give me all the electronic music bands having hometown in France or in Italy.	UNION	.020	.043	.027
176	What are the birthplace and the death place of Elvis Presley?	UNION	1	.500	.667
186	What are the birth names of Tom Cruise and Whoopi Goldberg?	UNION	1	.500	.667
190	What are the capitals of Italy and France?	UNION	1	.500	.667
201	How many children does Barack Obama have?	COUNT	1	1	1
231	Who are the author and publisher of Don Quixote?	UNION	1	.500	.667
239	Give me the population density of Warsaw as well as its population total.	UNION	1	.500	.667

Table 6

List of the questions answered by each system over the new queries added in MQALD.

dling this kind of questions, not only from an architectural point of view but also in creating and updating the

existing resources to allow these systems to be fairly compared.

This work aimed to give more insight about this specific issue related to Question Answering over Knowledge Graphs. This research area is fastly growing, and the availability of novel resources can significantly help develop novel solutions.

As future work, we plan to further encourage the research and development of new solutions to improve QA systems' performance over KG handling questions involving modifiers, which are very common in natural language. The first step towards this direction would be to revise the available resources properly, fix the datasets so that question/query pairs are compliant with the last version of DBpedia, and make sure that using a specific modifier is the only way to retrieve the right answers. Next, it would be beneficial for the whole community to expand this dataset further so that a sufficient number of questions properly represent all the modifiers available in SPARQL.

## Acknowledgement

This work was supported by the PRIN 2009 project: "Modelli per la personalizzazione di percorsi formativi in un sistema di gestione dell'apprendimento".

## References

- [1] E. Kaufmann and A. Bernstein, How Useful Are Natural Language Interfaces to the Semantic Web for Casual End-Users?, in: *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11-15, 2007*, K. Aberer, K. Choi, N.F. Noy, D. Allemang, K. Lee, L.J.B. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber and P. Cudré-Mauroux, eds, Lecture Notes in Computer Science, Vol. 4825, Springer, 2007, pp. 281–294. doi:10.1007/978-3-540-76298-0\_21.
- [2] F. Ozcan, A. Quamar, J. Sen, C. Lei and V. Efthymiou, State of the Art and Open Challenges in Natural Language Interfaces to Data, in: *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*, D. Maier, R. Pottinger, A. Doan, W. Tan, A. Alawini and H.Q. Ngo, eds, ACM, 2020, pp. 2629–2636. doi:10.1145/3318464.3383128.
- [3] A. Mishra and S.K. Jain, A survey on question answering systems with classification, *J. King Saud Univ. Comput. Inf. Sci.* **28**(3) (2016), 345–361. doi:10.1016/j.jksuci.2014.10.007.
- [4] I. Androutsopoulos, G.D. Ritchie and P. Thanisch, Natural language interfaces to databases - an introduction, *Nat. Lang. Eng.* **1**(1) (1995), 29–81. doi:10.1017/S135132490000005X.
- [5] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak and Z.G. Ives, DBpedia: A Nucleus for a Web of Open Data, in: *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007, Busan, Korea, November 11-15, 2007*, K. Aberer, K. Choi, N.F. Noy, D. Allemang, K. Lee, L.J.B. Nixon, J. Golbeck, P. Mika, D. Maynard, R. Mizoguchi, G. Schreiber and P. Cudré-Mauroux, eds, Lecture Notes in Computer Science, Vol. 4825, Springer, 2007, pp. 722–735. doi:10.1007/978-3-540-76298-0\_52.
- [6] D. Vrandečić and M. Krötzsch, Wikidata: a free collaborative knowledgebase, *Commun. ACM* **57**(10) (2014), 78–85. doi:10.1145/2629489.
- [7] 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) (2017), 895–920. doi:10.3233/SW-160247.
- [8] 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) (2018), 529–569. doi:10.1007/s10115-017-1100-y.
- [9] 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: *Joint Proceedings of BLINK2017: 2nd International Workshop on Benchmarking Linked Data and NLIWoD3: Natural Language Interfaces for the Web of Data co-located with 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria, October 21st - to - 22nd, 2017*, R. Usbeck, A.N. Ngomo, J. Kim, K. Choi, P. Cimiano, I. Fundulaki and A. Krithara, eds, CEUR Workshop Proceedings, Vol. 1932, CEUR-WS.org, 2017. http://ceur-ws.org/Vol-1932/paper-02.pdf.
- [10] M. Dubey, S. Dasgupta, A. Sharma, K. Höffner and J. Lehmann, AskNow: A Framework for Natural Language Query Formalization in SPARQL, in: *The Semantic Web. Latest Advances and New Domains - 13th International Conference, ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Proceedings*, H. Sack, E. Blomqvist, M. d'Aquin, C. Ghidini, S.P. Ponzetto and C. Lange, eds, Lecture Notes in Computer Science, Vol. 9678, Springer, 2016, pp. 300–316. doi:10.1007/978-3-319-34129-3\_19.
- [11] 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). http://arxiv.org/abs/1907.09361.
- [12] L. Siciliani, D. Diefenbach, P. Maret, P. Basile and P. Lops, Handling Modifiers in Question Answering over Knowledge Graphs, in: *AI\*IA 2019 - Advances in Artificial Intelligence - XVIIIth International Conference of the Italian Association for Artificial Intelligence, Rende, Italy, November 19-22, 2019, Proceedings*, M. Alviano, G. Greco and F. Scarcello, eds, Lecture Notes in Computer Science, Vol. 11946, Springer, 2019, pp. 210–222. doi:10.1007/978-3-030-35166-3\_15.
- [13] K.D. Bollacker, C. Evans, P. Paritosh, T. Sturge and J. Taylor, Freebase: a collaboratively created graph database for structuring human knowledge, in: *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2008, Vancouver, BC, Canada, June 10-12, 2008*, J.T. Wang, ed., ACM, 2008, pp. 1247–1250. doi:10.1145/1376616.1376746.
- [14] 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 Lan-*

- 1 *guage Processing, EMNLP 2013, 18-21 October 2013, Grand*  
 2 *Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT,*  
 3 *a Special Interest Group of the ACL, ACL, 2013, pp. 1533–*  
 4 *1544. <https://www.aclweb.org/anthology/D13-1160/>.*
- [15] A. Bordes, N. Usunier, S. Chopra and J. Weston, Large-scale Simple Question Answering with Memory Networks, *CoRR abs/1506.02075* (2015). <http://arxiv.org/abs/1506.02075>.
- [16] P. Trivedi, G. Maheshwari, M. Dubey and J. Lehmann, LC-QuAD: A Corpus for Complex Question Answering over Knowledge Graphs, in: *The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part II*, C. d’Amato, M. Fernández, V.A.M. Tamma, F. Lécué, P. Cudré-Mauroux, J.F. Sequeda, C. Lange and J. Heflin, eds, Lecture Notes in Computer Science, Vol. 10588, Springer, 2017, pp. 210–218. doi:10.1007/978-3-319-68204-4\_22.
- [17] 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: *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Auckland, New Zealand, October 26-30, 2019, Proceedings, Part II*, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I.F. Cruz, A. Hogan, J. Song, M. Lefrançois and F. Gandon, eds, Lecture Notes in Computer Science, Vol. 11779, Springer, 2019, pp. 69–78. doi:10.1007/978-3-030-30796-7\_5.
- [18] Z. Jia, A. Abujabal, R.S. Roy, J. Strötgen and G. Weikum, TempQuestions: A Benchmark for Temporal Question Answering, in: *Companion of the The Web Conference 2018 on The Web Conference 2018, WWW 2018, Lyon, France, April 23-27, 2018*, P. Champin, F. Gandon, M. Lalmas and P.G. Ipeirotis, eds, ACM, 2018, pp. 1057–1062. doi:10.1145/3184558.3191536.
- [19] A. Talmor and J. Berant, The Web as a Knowledge-Base for Answering Complex Questions, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, M.A. Walker, H. Ji and A. Stent, eds, Association for Computational Linguistics, 2018, pp. 641–651. doi:10.18653/v1/n18-1059.
- [20] T.S. Costa, S. Gottschalk and E. Demidova, EventQA: A Dataset for Event-Centric Question Answering over Knowledge Graphs (2020), 3157–3164. doi:10.1145/3340531.3412760.
- [21] R. Usbeck, M. Röder, A.N. Ngomo, C. Baron, A. Both, M. Brümmer, D. Ceccarelli, M. Cornolti, D. Cherix, B. Eickmann, P. Ferragina, C. Lemke, A. Moro, R. Navigli, F. Piccinno, G. Rizzo, H. Sack, R. Speck, R. Troncy, J. Waitelonis and L. Wesemann, GERBIL: General Entity Annotator Benchmarking Framework, in: *Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015*, A. Gangemi, S. Leonardi and A. Panconesi, eds, ACM, 2015, pp. 1133–1143. doi:10.1145/2736277.2741626.
- [22] 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) (2019), 293–304. doi:10.3233/SW-180312.
- [23] 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: *Semantic Web Challenges - 4th SemWebEval Challenge at ESWC 2017, Portoroz, Slovenia, May 28 - June 1, 2017, Revised Selected Papers*, M. Dragoni, M. Solanki and E. Blomqvist, eds, Communications in Computer and Information Science, Vol. 769, Springer, 2017, pp. 59–69. doi:10.1007/978-3-319-69146-6\_6.
- [24] R. Usbeck, A.N. Ngomo, F. Conrads, M. Röder and G. Napolitano, 8th Challenge on Question Answering over Linked Data (QALD-8) (invited paper) **2241** (2018), 51–57. <http://ceur-ws.org/Vol-2241/paper-05.pdf>.
- [25] R. Usbeck, R.H. Gusmita, A.N. Ngomo and M. Saleem, 9th Challenge on Question Answering over Linked Data (QALD-9) (invited paper), in: *Joint proceedings of the 4th Workshop on Semantic Deep Learning (SemDeep-4) and NLIWoD4: Natural Language Interfaces for the Web of Data (NLIWOD-4) and 9th Question Answering over Linked Data challenge (QALD-9) co-located with 17th International Semantic Web Conference (ISWC 2018), Monterey, California, United States of America, October 8th - 9th, 2018*, K. Choi, L.E. Anke, T. Declerck, D. Gromann, J. Kim, A.N. Ngomo, M. Saleem and R. Usbeck, eds, CEUR Workshop Proceedings, Vol. 2241, CEUR-WS.org, 2018, pp. 58–64. <http://ceur-ws.org/Vol-2241/paper-06.pdf>.
- [26] S. Hu, L. Zou and X. Zhang, A State-transition Framework to Answer Complex Questions over Knowledge Base, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, E. Riloff, D. Chiang, J. Hockenmaier and J. Tsujii, eds, Association for Computational Linguistics, 2018, pp. 2098–2108. doi:10.18653/v1/d18-1234.
- [27] D. Diefenbach, A. Both, K. Singh and P. Maret, Towards a question answering system over the Semantic Web, *Semantic Web* **11**(3) (2020), 421–439. doi:10.3233/SW-190343.
- [28] D. Diefenbach, K.D. Singh and P. Maret, WDAqua-core0: A Question Answering Component for the Research Community, in: *Semantic Web Challenges - 4th SemWebEval Challenge at ESWC 2017, Portoroz, Slovenia, May 28 - June 1, 2017, Revised Selected Papers*, M. Dragoni, M. Solanki and E. Blomqvist, eds, Communications in Computer and Information Science, Vol. 769, Springer, 2017, pp. 84–89. doi:10.1007/978-3-319-69146-6\_8.
- [29] D. Vollmers, R. Jalota, D. Moussallem, H. Topiwala, A.N. Ngomo and R. Usbeck, Knowledge Graph Question Answering using Graph-Pattern Isomorphism, *CoRR abs/2103.06752* (2021). <https://arxiv.org/abs/2103.06752>.