8th Challenge on Question Answering over Linked Data (QALD-8)

Ricardo Usbeck[1], Axel-Cyrille Ngonga Ngomo[1], Felix Conrads[1], Michael Röder[1], and Giulio Napolitano[2]

Data Science Group, Paderborn University, Germany
ricardo.usbeck@uni-paderborn.de
Fraunhofer-Institute IAIS, Sankt Augustin, Germany
giulio.napolitano@iais.fraunhofer.de

1 Introduction

Recent years have seen a growing amount of research on question answering (QA) over Semantic Web data, shaping an interaction paradigm that allows end users to profit from the expressive power of Semantic Web standards. At the same time, this system hides their complexity behind an intuitive and easy-to-use interface. However, the growing amount of data has led to a heterogeneous data landscape where QA systems struggle to keep up with the volume, variety and veracity of the underlying knowledge.

The Question Answering over Linked Data (QALD) challenge aims to provide an up-to-date benchmark for assessing and comparing state-of-the-art-systems that mediate between a user, expressing his or her information need in natural language, and RDF data. In the past few years, more than **38 research groups took part in the last eight QALD challenges**. The QALD challenge targets all researchers and practitioners working on querying Linked Data, natural language processing for question answering, multilingual information retrieval and related topics. The main goal is to gain insights into the strengths and shortcomings of different approaches and into possible solutions for coping with the large, heterogeneous and distributed nature of Semantic Web data.

QALD[1] has a 7-year history of developing a benchmark that is increasingly being used as a standard evaluation venue for question answering over Linked Data. Overviews of past instantiations of the challenge are available from the CLEF Working Notes as well as ESWC proceedings, see Table 1.

This article will give a technical overview of the task and results of the 8th Question Answering over Linked Data challenge.

2 Dataset and Task

The key challenge for QA over Linked Data is to translate a user's information need into such a form that it can be evaluated using standard Semantic Web

[1] http://www.sc.cit-ec.uni-bielefeld.de/qald/

QALD	URI to proceedings
7	https://link.springer.com/chapter/10.1007/978-3-319-69146-6_6
6	https://www.springer.com/us/book/9783319465647
5	https://ceur-ws.org/Vol-1391/173-CR.pdf
4	https://ceur-ws.org/Vol-1180/CLEF2014wn-QA-UngerEt2014.pdf
3	https://pub.uni-bielefeld.de/download/2685575/2698020
2	https://ceur-ws.org/Vol-913/
1	https://qald.sebastianwalter.org/1/documents/qald-1-proceedings.pdf

Table 1. QALD proceedings and their URIs.

query processing and inferencing techniques. The main task of QALD therefore is the following:

> Given one or several RDF dataset(s) as well as additional knowledge sources and natural language questions or keywords, return the correct answers or a SPARQL query that retrieves these answers.

Data format

All the data for the tasks can be found in our project repository `https://github.com/ag-sc/QALD/tree/master/8/data`. We encouraged the use of QALD-JSON format[2] as communication format between the systems and the GERBIL QA respectively HOBBIT platform:

```
{ "id":"3",
  "answertype":"resource",
  "aggregation":false,
  "onlydbo":true,
  "hybrid":false,
  "question":[
    {
      "language":"en",
      "string":"Who was the wife of U.S. president
         Lincoln?",
      "keywords":"U.S. president, Lincoln, wife"
    },
    {
      "language":"nl",
      "string":"Wie was de vrouw van de Amerikaanse
         president Lincoln?",
      "keywords":"vrouw, president van America, Lincoln
         "
```

[2] `https://github.com/AKSW/gerbil/wiki/Question-Answering` and the results are formatted according to `https://www.w3.org/TR/sparql11-results-json/`

```
    }
  ],
  "query":{
    "sparql":"PREFIX dbo:<http://dbpedia.org/ontology/>
        PREFIX res:<http://dbpedia.org/resource/>
        SELECT DISTINCT ?uri
        WHERE {res:Abraham_Lincoln dbo:spouse ?uri.}"
  },
  "answers":[
    {
      "head":{
        "vars":[
          "uri"
        ]
      },
      "results":{
        "bindings":[
          {
            "uri":{
              "type":"uri",
              "value":"http://dbpedia.org/resource/
                 Mary_Todd_Lincoln"
            }
          }
        ]
    }}]}
```


Task: Multilingual question answering over DBpedia

Given the diversity of languages used on the web, there is an increasing need to facilitate multilingual access to semantic data. The core task of QALD is thus to retrieve answers from an RDF data repository given an information need expressed in a variety of natural languages.

Training data. The underlying RDF dataset was DBpedia 2016-10. The training data consists of 219 questions compiled and curated from previous challenges. The questions are be available in eight different languages (English, Spanish, German, Italian, French, Dutch, Romanian and Farsi). Those questions are general, open-domain factual questions, for example:

(en) *Which book has the most pages?*
(de) *Welches Buch hat die meisten Seiten?*
(es) *¿Que libro tiene el mayor numero de paginas?*
(it) *Quale libro ha il maggior numero di pagine?*
(fr) *Quel livre a le plus de pages?*

(nl) *Welk boek heeft de meeste pagina's?*
(ro) *Ce carte are cele mai multe pagini?*

The questions vary with respect to their complexity, including questions with counts (e.g., *How many children does Eddie Murphy have?...*), superlatives (e.g., *Which museum in New York has the most visitors?*), comparatives (e.g., *Is Lake Baikal bigger than the Great Bear Lake?*), and temporal aggregators (e.g., *How many companies were founded in the same year as Google?*). Each question is annotated with a manually specified SPARQL query and answers. In the above case, the SPARQL query looks as follows:

```
SELECT DISTINCT ?uri
WHERE {
    ?uri a <http://dbpedia.org/ontology/Book> .
    ?uri <http://dbpedia.org/ontology/numberOfPages> ?n .
}
ORDER BY DESC(?n)
OFFSET 0 LIMIT 1
```

And the answer is `<http://dbpedia.org/resource/The_Tolkien_Reader>`.

Test Data. The test dataset consists of 42 similar manually created questions. However, this year we decided to increase the complexity of the test data and add several other question types including questions according to RDF types (e.g., *What is backgammon?...*) or questions demanding mathematical operations (e.g., *What is the radius of the earth?...*). They are compiled from existing, real-world question and query logs, in order to provide unbiased questions expressing real-world information needs. The questions were manually curated to ensure a high quality standard.

Evaluation Metric. The 8th QALD challenge provides an automatic evaluation tool, namely GERBIL QA [3][3], which is open source and available for everyone to re-use. The GERBIL QA platform is accessible online, so that participants can simply upload the answers produced by their system or even check their system via a webservice. Each experiment has a citable, time-stable and archivable URI that is both human- and machine-readable. However, participating systems have to provide a webservice to participate in the final challenge.[4]

The QA systems were evaluated with respect to precision and recall:

$$\text{recall}(q) = \frac{\text{number of correct system answers for } q}{\text{number of gold standard answers for } q}$$

$$\text{precision}(q) = \frac{\text{number of correct system answers for } q}{\text{number of system answers for } q}$$

The metrics use the following additional semantic information:

[3] http://gerbil-qa.aksw.org/gerbil/
[4] https://github.com/dice-group/gerbil/wiki/Question-Answering#web-service-interface

- If the golden answerset is empty and the system does respond with an empty answer, we set precision, recall and F-measure to 1.
- If the golden answerset is empty but the system responds with any answerset, we set precision, recall and F-measure to 0.
- If there is a golden answer but the the QA system responds with an empty answerset, we assume the system could not respond. Thus we set the precision to 0 and the recall and F-measure to 0.
- In any other case, we calculate the standard precision, recall and F-measure per question.

The evaluation also computed the macro and micro F-measure of a system over all test questions. That is, for micro F-measure, we summed up all true and false positives and negatives and calculated the precision, recall and F-measure at the end, while for the macro measures, we calculated precision, recall and F-measure per question and averaged the values at the end.

For the final evaluation, we focused only on the Macro F1 QALD metric. That is, we decided to have a metric more comparable to older QALD challenges and also to follow community requests.[5] This metric uses the previously mentioned additional semantic information with the following exception:

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

3 Participating systems

After eight registrations, only two teams were able to join the final evaluation. We added QAKIS [1] as a baseline.

WDAqua [2] is a rule-based system using a combinatorial approach to generate SPARQL queries from natural language questions, leveraging the semantics encoded in the underlying knowledge base. It can answer questions on both DBpedia (supporting English) and Wikidata (supporting English, French, German and Italian). The system, which does not require training, participated in Tasks 1 and 4 of the challenge.

ganswer2 [4] has participated outside the actual challenge this year as a system without a paper submission in Task 1. Zou et al. use a graph-based approach to generate a semantic query graph, which reduced the transformation of natural language to SPARQL to a subgraph matching problem.

4 Results

All QA systems were run on the QALD-8 train and test dataset in English and GERBIL QA version 0.2.1 and you can find FAIR experiment data at `http://w3id.org/gerbil/qa/experiment?id=201710220000`.

[5] `https://github.com/dice-group/gerbil/issues/211`

QALD-8 test introduced some **curve balls**.[6] Since we used real world queries extracted from logs of different search engines, the queries are dirty, i.e. contain misspellings, have new forms, (e.g., "Who is X?", asking for a description) or have only weakly defined answer types, i.e., sets of strings instead of URI sets. Note also that the numbers reported here may differ from the publications of the participants, as these figures were not available at the time of participant paper submission.

Since both real systems, gAnswer and WDAqua-core0 were so close together and the Macro F1 QALD metric was new, the organizing committee gave a tied result and both systems won the QALD-8 challenge.

Annotator	Dataset	Macro F1	Macro Precision	Macro Recall	avg millis/doc	Macro F1 QALD
QAKIS	Test	0.0563	0.061	0.0528	15,414	0.0566
gAnswer	Test	0.388	0.3862	0.3902	1,919	0.3882
WDAqua-core0	Test	0.3872	0.3912	0.4065	1,725	0.3987

Table 2. Final results on the test dataset of QALD-8.

5 Summary

The QALD-8 challenge focused on the successful and long running multilingual QA task. For the first time, the participating teams were required to provide webservices of their systems to participate in the challenge, which will in turn support comparable research in the future. In this challenge, we also changed the underlying evaluation platform to account for the need for comparable experiments via webservices in contrast to former XML/JSON file submissions. This increased the entrance requirements for participating teams but ensures long term comparability of the system performance and a fair and open challenge.

In the future, we will further simplify the participation process and offer leaderboards prior to the actual challenge to allow participants to see their performance beforehand. After feedback from the authors, we will likely add new key performance indicators for the capability of a system to know which questions it cannot answer and take confidence scores for answers into account. Moreover, we will remove most of the curve ball questions to reflect the original character of the QALD challenge, which provides a clean and linguistically challenging benchmark.

Acknowledgments. This work has been supported by the Federal Ministry of Transport and Digital Infrastructure (BMVI) in the LIMBO research project

[6] `https://github.com/ag-sc/QALD/blob/master/8/data/qald-8-test-multilingual.json`

(grant no. 19F2029C) and the in the OPAL research project (grant no. 19F2028A) as well as by the German Federal Ministry of Education and Research (BMBF) within 'KMU-innovativ: Forschung für die zivile Sicherheit' in particular 'Forschung für die zivile Sicherheit' and the project SOLIDE (no. 13N14456).

The authors gratefully acknowledge financial support from the German Federal Ministry of Education and Research within Eurostars, a joint programme of EUREKA and the European Community under the project E! 9367 DIESEL and E! 9725 QAMEL.

We also want to thank Christina Unger and Sebastian Walter for supporting this challenge.

Data experts for creating QALD-8:

– Rricha Jalota
Paderborn University, Germany
– Paramjot Kauer
Paderborn University, Germany
– Abdullah Ahmed
Paderborn University, Germany
– Felix Conrads
Paderborn University, Germany and InfAI e.V. Leipzig

References

1. Elena Cabrio, Julien Cojan, Fabien Gandon, and Amine Hallili. Querying Multilingual DBpedia with QAKiS. In *ESWC*, pages 194–198, 2013.
2. Dennis Diefenbach, Kamal Deep Singh, and Pierre Maret. Wdaqua-core0: A question answering component for the research community. In Mauro Dragoni, Monika Solanki, and Eva Blomqvist, editors, *Semantic Web Challenges - 4th SemWebEval Challenge at ESWC 2017, Portoroz, Slovenia, May 28 - June 1, 2017, Revised Selected Papers*, volume 769 of *Communications in Computer and Information Science*, pages 84–89. Springer, 2017.
3. Ricardo Usbeck, Michael Röder, Michael Hoffmann, Felix Conrad, Jonathan Huthmann, Axel-Cyrille Ngonga-Ngomo, Christian Demmler, and Christina Unger. Benchmarking question answering systems. *Semantic Web Journal*, 2018.
4. Lei Zou, Ruizhe Huang, Haixun Wang, Jeffrey Xu Yu, Wenqiang He, and Dongyan Zhao. Natural language question answering over rdf: a graph data driven approach. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 313–324. ACM, 2014.