Challenges of Entity Linking in KGQA datasets

Kanchan Shivashankar, Nadine Steinmetz

Bergische Universität Wuppertal, Wuppertal, Germany

Abstract

Entity Linking (EL) is the process of extracting surface form entities from the text and linking it to the knowledge graph ontology. It is one of the challenges facing Question Answering on Knowledge Graph (KGQA). Using correct entities in the SPARQL query, determines the outcome of KGQA systems. Ongoing research provides solutions across neural and analytical approaches. However, lexical gap and ambiguity are the two biggest challenges in EL. The study of different entity forms in questions can benefit the outcome of EL systems. In this paper, we analyse the existing gap between natural language entity forms and three different knowledge graph ontologies (DBpedia, Wikidata and DBLP-RDF). The closeness in match between the surface entity forms and respective knowledge graph is classified and compared. We present our findings on four benchmark KGQA datasets (QALD 9, LC-QuAD 2.0, SMART 2022 EL and DBLP-QuAD) and recommend possible solutions for efficient EL.

Keywords

Entity linking, KGQA

1. Introduction

Question Answering on Knowledge Graphs (KGQA) is a fact-based discipline and aims to transform Natural Language (NL) question to a SPARQL query to be answered on an RDF graph. Important substeps to generate the SPARQL query are the identification of named entities and the respective relations within the NL. Both are essential to create the correct RDF triples for the SPARQL query.

In general, Entity Linking (EL) deals with the lexical gap and disambiguation problems. The lexical gap refers to the problem of missing references from a surface form to the correct named entity due to incomplete dictionaries. Disambiguation problems are a general problem for NL texts, but for QA scenarios more serious due to little context information in the question. In addition to these two challenges, KGQA datasets are sometimes set up with extra hard entity linking tasks. With this paper, we discuss the complexity of linking surface forms in NL questions to the respective named entities of the relevant knowledge graph. In order to identify specific challenges, we analyzed four different datasets for KGQA and specifically EL for KGQA: QALD 9, LC-QuAD 2.0, SMART 2022 EL, and DBLP-QuAD. Thereby, we also include three different knowledge graphs in our analyses: Wikidata, DBpedia, and DBLP-RDF.

Besides some statistics about the datasets, we provide elaborate analyses on the different

NLP4KGC@SEMANTICS 2023 EU: 19th International Conference on Semantic Systems, September 20-22, 2023, Leipzig, Germany

shivashankar@uni-wuppertal.de (K. Shivashankar); steinmetz@uni-wuppertal.de (N. Steinmetz)

0000-0003-3601-7579 (N. Steinmetz)

challenges regarding linking named entities. Some of these challenges are common for all or most of the datasets, but we also take a look at challenges that are very specific for the dataset and/or the relevant knowledge graph.

Thereby, we provide an overview of characteristics of the datasets regarding entity linking and also an extensive discussion on how to overcome those challenges. Our contribution enables KGQA developers to skip detailed researches on the entity linking task regarding the respective dataset and fully focus on the development of the QA application.

The remainder of the paper is structured as follows: the next section introduces some related work on the aspect of KGQA and dataset overviews. Section 3 introduces the datasets under consideration and provides some first-hand statistics. Section 4 describes our findings regarding challenges on entity linking and classifies the identified issues. Section 5 summarizes our findings and the overview of KGQA datasets.

2. Related Work

Entity linking is one of the important subtasks in KGQA. The two main challenges of EL are attributed to different variations of entity representation in natural language and entity ambiguity.

Entity linking consists of three modules: candidate entity generation, candidate entity ranking and unlinkable mention prediction [1]. Candidate generation results in a list of candidates for the entity mentioned in the text by filtering out irrelevant ones, which are then ranked based on relevance to the context of the mention. Ideally, the knowledge base should define all the entities, but in reality this is not always true. In case an entity mention does not have a corresponding link in the knowledge base, the unlinkable problem of entity mentions should be handled.

KGQA problems can be solved with two different types of systems. Recent approaches in KGQA focus on neural approaches that include an end-to-end system and are treated as a machine translation problems. The Neural Machine Translation (NMT) approach translates natural language questions to structured SPARQL queries [2]. An NMT system is only as good as the training data and can have problems with EL if it encounters out-of-vocabulary entities in the test data. It provides no transparency and little to no control over how EL is performed. A better way is to approach this as a pipeline with separate modules to handle different sub-tasks such as entity linking, relation linking, and query formulation [3]. To this accord, NMT when combined with EL and RL modules has shown to improve the results compared to pure NMT approaches [4]. The transformer based NMT model with EL achieves the best result of 67.9% accuracy and the pure NMT model only achieves 50% accuracy for the QALD 9 dataset.

Understanding the importance of EL, benchmark datasets have emerged for evaluating state-of-the-art systems in EL for KGQA. The SMART 2022 challenge for EL and RL released training and test dataset for the DBpedia and Wikidata knowledge graphs. CLOCQ is an unsupervised framework that computes top-k linking for each mention in the question and then adapted to prune irrelevant mentions [5]. The results indicate a trade-off between precision and recall; higher k value boosts recall while lower values of k leads to high precision. On the other hand, [6] uses KG refinement techniques to solve EL. The question is used to generate candidate terms, then using the KG as an external knowledge the candidate terms are queried. However, the approach fails to identify longer entities that include stopwords in them.

Other popular models for EL are FALCON [7], a rule based approach for EL and RL for DBpedia and BLINK by Facebook AI Research [8], which is a BERT based tool for EL. The latter is a very efficient EL tool with state-of-the-art performance in zero shot (without external knowledge) setup. In [9], entity linking is proposed using a combination of trained and analytical approaches using Abstract Meaning Representation (AMR). The AMR graph generation model is trained on augmented data, using which surface form of entities are identified and mapped to the knowledge graph using an entity dictionary.

Part of the EL challenge is also dependent on the structure of the knowledge graph. If the attributes of the KG vary, the results of EL may be impacted. For example, Wikidata contains seventy-three million items that can be interpreted as entities due to the existence of an 'is instance' property, in comparison, DBpedia contains around 5 million entities. The 'is instance' property includes a much broader scope of entities than the ones interpreted as entities for DBpedia [10]. Additionally, the information in Wikidata is maintained by an active community, which leads to frequent updates, compared to DBpedia. These updates make it harder to replicate research results.

In the next section, we take a look at the datasets under consideration and discuss their attributes and pre-processing for our analyses.

3. Datasets

3.1. Overview

The findings are presented for four different KGQA datasets based on three different knowledge graphs for English language questions. The different KGs and datasets used for our analyses are described below.

1. DBpedia Knowledge Graph
    a) *QALD 9 (Question Answering over Linked Data)* [11]
    The 9th challenge on Question Answering over Linked Data released their dataset in 2018. It contains more heterogeneous questions, such as, questions with counts, superlatives, comparatives and temporal aggregators compared to previous QALD challenges. Each question is annotated with manually specified SPARQL queries. For our analyses, we use the training and test data.
    b) *SMART 2022 Entity Linking: DBpedia* [1]
    The SMART 3.0 challenge of ISWC 2022 released a specific dataset for the EL task on the DBpedia graph. The dataset provides questions and respective entity lists. We consider only the training dataset for our analyses as the entity lists for the test data is not published.
2. Wikidata Knowledge Graph
    a) *LC-QuAD 2.0 - Wikidata*[12]
    LC-QuAD(Large-scale Complex Question Answering Dataset) is the biggest dataset

[1]https://smart-task.github.io/2022/

Table 1
KGQA dataset statistics for EL.

Datasets	No. of questions	Questions with NE	Total no. of entities
QALD 9 Train	408	366	414
QALD 9 Test	150	125	135
LC-QuAD 2.0 Train	24180	23596	32561
LC-QuAD 2.0 Test	6046	5890	8175
SMART EL Train	49976	49976	56629
DBLP-QuAD Train	7000	7000	9407
DBLP-QuAD Valid	1000	1000	1357
DBLP-QuAD Test	2000	2000	2738

of complex questions with corresponding SPARQL queries available for KGQA. It is compatible with Wikidata and DBpedia 2018 knowledge graphs.

b) *SMART 2022 Entity Linking: Wikidata*
As described above for DBpedia, the challenge also released a set of questions used for identifying entities in Wikidata knowledge graph.

3. DBLP Knowledge Graph
   a) *DBLP-QuAD*[2]
   DBLP is an online reference for bibliographic information on major computer science publications. It contains over 4.4 million publications by more than 2.2 million authors. The DBLP-QuAD dataset consists of 10,000 question-answer pairs with SPARQL queries, that is split into train, valid and test data. The SPARQL queries can be executed over an DBLP-RDF KG.

The different datasets provide a more diverse sample of questions and thus a variety of entity form representations. Entity linking is also dependent on the structure of the underlying KG and can also be an influencing factor in our findings. The table 1 provides basic statistics on the different datasets. The first column indicates the total number of questions present in the respective dataset, whereas, the second column provides the number of questions out of the total that requires entity linking. The third column of values are the total number of entities present in the dataset: a higher value indicates presence of questions with multiple entities.

3.2. Preparation

The named entities required to answer the question is not a prior knowledge in KGQA. The data needs to be pre-processed, in order to extract entities from the question. This involves two steps,

1. Fetch the list of named entities from the SPARQL query for each question.
2. Match the list of labels from (1) to entities mentioned in the question.

For the first step, the entity URI (Uniform Resource Identifier) for some of the datasets (DBLP and SMART 2022 EL) are already available for each record in the dataset and must not be

[2] https://doi.org/10.5281/zenodo.7643971

Table 2
Entity match distribution for KGQA datasets.

KG	DBpedia			Wikidata		DBLP		
Datasets	QALD 9		SMART 2022 EL	LC-QuAD 2.0		DBLP-QuAD		
	Train	Test	Train	Train	Test	Train	Valid	Test
Total Entities	414	132	56629	56629	32561	8175	9407	1357	2738
Exact Match	314	98	42880	45086	24929	6200	5174	740	1454
Partial Match	80	23	8813	10036	6099	1596	4002	577	1184
No Match	4	3	959	1269	1344	339	9	1	0
URI Mismatch	16	8	3977	238	189	40	222	39	100

extracted from the SPARQL query. For the other datasets (QALD 9, LC-QuAD 2.0), we extracted the URI from the given SPARQL query.

The URI of the entity is a unique identifier in the KG, but we need a NL label for the entity to compare with the question text. We therefore collect a list of labels for each URI. For this, we use entity dictionaries for each KG. The entity form in NL can be represented in various forms. To compensate for this challenge, we include list of labels that also contain variations of the original label to find better matches in the question. For DBpedia, the entity dictionary consists of the RDF label, the redirect label and respective disambiguation labels. For Wikidata, we utilize the main and alternative labels provided in the KG. DBLP-RDF mainly only provides main labels for authors and publications. Thereby, the entity dictionaries are used reversely to collect the labels for the entity URIs of the question/query.

For the example question "What is the time zone of Salt Lake City?", the URI for the entity *Salt Lake City* in the DBpedia is 'http://dbpedia.org/resource/Salt_Lake_City'. Our dictionary contains the following labels for the entity: salt lake city, saltlake city, salt lake cit, salt lake city ut, slc, salt lake ut, slc ut.

This concludes the data preparation process. In the next step, the entity labels are matched against the question text to find complete, partial or no matches. Thus, providing the necessary information to perform an in-depth analysis, which is discussed in the next section.

4. Findings

The datasets have been analyzed to identify the representation of named entities in questions. We discuss in detail the different forms of entity representation and its challenges to the task of EL.

4.1. Overview

As explained in the previous section, additional data is required i.e, list of entity labels for our analysis. We tried to identify the list of entity URIs from the dataset, directly (entity list provided) or indirectly (SPARQL query). With the help of the URI, we then collect labels for the entity and find a match to keywords from the question. This is the entity form representation in NL.

The difficulty of performing EL on the questions depends on the closeness measure between the two terms (the keywords in the question and labels defined in the KG) also referred to as the lexical gap. For this closeness measure, we identified three groups: (i) Exact Match (ii) Partial Match and (iii) No Match. In addition, we discuss EL issues caused due to the dataset and KG dependencies. Each category has a specific set of issues which are grouped into different classes and discussed in the next subsection. Table 2 illustrates the distribution of the entity representation in questions across different categories for each dataset.

4.2. Entity Representations

4.2.1. Exact Match

The simplest case, when the entity labels in the question and knowledge graphs match. For instance, the examples below contain the full label of the required entity in the question:

qald9_train_7: [3] Where did <u>Abraham Lincoln</u> die?

matched_entity: abraham lincoln

dblp_train_Q0004: What is the webpage of <u>Ravi Kumar</u>?

matched_entity: ravi kumar

smart2022_train_wikidata_4: What is <u>anthony dowell</u>'s gender?

matched_entity: anthony dowell

This is the expectation in an ideal case. However, due to the complex nature of NL, questions can be formed in different ways that mean the same and thus entities in questions can also be referred to in more than one way. The next two cases demonstrate this ambiguity in depth.

4.2.2. Partial Match

If the complete entity label is not present in the question, the exact match condition fails and we resort to finding a partial match for the entity. Partial match checks if parts of the entity label or its variation is present in the question. The examples below discuss the different forms of entity representation in the question. For instance, with partial names, referring to a popular baseball team *Boston Red Sox* as *Red Sox*, using the adjective form *Swiss* to refer to the country (noun) *Switzerland*, using abbreviations *ESA* for *European Space Agency* etc.

qald9_train_45: Where do the <u>Red Sox</u> play?

matched_entity_pair: (Boston Red Sox, Red Sox)

[3] qald9_train_7: The question belongs to qald9 dataset, training data with id=7

qald9_train_42: Give me all Swiss non-profit organizations.

matched_entity_pair: (Swisserland, Swiss)

qald9_train_73: Give me all ESA astronauts.

matched_entity_pair: (European Space Agency, ESA)

dblp_train_Q0024: What are the papers written by Jialu D.?

matched_entity_pair: (du, jialu, jialu)

smart22_train_dbp_49838: Descartes was influenced by which polymath?

matched_entity_pair: (rené descartes, descartes)

In partial match cases, we make use of the list of labels created for each entity URI. This list contains variations to compensate for a lexical gap in entity representation, context words and also in some cases spelling mistakes attributed to human error. But, as discussed above, these dictionaries are far from complete in many cases.

4.2.3. No Match

If the first two cases i.e, exact match and partial match fails, it is likely that the entity word is absent in the question or is implied through context or some relation. This is one of the challenges in the EL task, to link an entity to a question that is not explicitly mentioned, but implied. We look at various examples in this section which has been grouped into different classes.

1. Entity absence: If the entity word is not explicitly mentioned in the question.

   smart2022_train_wikidata_49912: What coast does was nominated for Classical Album of the Year declare nationality?

   expected_entity: Marc-Éric Gueï

   lcquad_test_wikidata_49912: What medication helps you urinate?

   expected_entity: water

2. Ambiguous entity: The entity mentioned in the question is unclear and cannot be matched to a specific entity in the KG.

   dblp_train_Q0373 Name the venue in which D. published the paper about Sequence spread.

expected_entity: Brian D. Woerner

smart2022_train_dbp_48861 Under which label does stephjones2 work?

expected_entity: steph jones

3. Region or Language specific: The entity is linked to another word that is specific to a region or language. For example the word 'cosmonauts' refers to soviet or russian space travelers and thus the entities expected need to be inferred through the word 'cosmonaut' or referring to the spanish name 'el bola' for the english version of the movie 'pellet'.

   qald9_train_41: Give me all cosmonauts.

   expected_entity: Soviet Union, Russia

   smart2022_train_dbp_48852: Which country did the film pellet originate in?

   expected_entity: el bola

4. Entities used as an adjective: Named entities (e.g. Denmark) represented in its adjective form (e.g. Danish).

   qald9_train_51: Give me all Danish movies.

   expected_entity: Denmark

   qald9_test_157: Give me English actors starring in Lovesick.

   expected_entity: United Kingdom

4.2.4. URI Mismatch

If entities exist in the question, the entity label is extracted from querying the URI in the respective KG. But what if the URI is invalid? In this case, we observed that a URI mismatch occurs. This is due to issues specific to dataset or KG structure that can be an impediment in efficient EL.

1. DBpedia: Entity URIs in DBpedia contains its natural language id in it: 'http://dbpedia.org/resource/Universal_Studios'. If the label is updated or changed: 'http://dbpedia.org/resource/Universal_Pictures' in the KG after the dataset was created. This can lead to incompatibility between the KG and the dataset leading to URI mismatch cases.

   qald9_train_15: Who is the owner of Universal Studios?

   expected_entity: Universal Pictures

   qald9_train_99: Who killed John Lennon?

   expected_entity: Death of John Lennon

qald9_train_115: In which U.S. state is Mount McKinley located?

expected_entity: Denali

2. Wikidata: Similar to DBpedia, some of the Wikidata entities with the id, are also updated or redirected to a new URI id, causing URI mismatch. Additionally, Wikidata also identifies objects with 'instance-of' relation(P31) as entities (for example: entity:*Uttarakhand* is an instance of property:'state of India'), which is considered as a property in DBpedia ontology.

   lcquad_train_21961: What writing system does Arabic use?

   expected_entity: alphabetic writing system (Q2182919 redirected from Q11077643)

   lcquad_train_17517: Which birth language is Chizoba Ejike?

   expected_entity: Q21197460 entity does not exist

   smart2022_train_wikidata_95: What is an example of a crime fiction?

   expected_entity: crime fiction (Q5937792 redirected from Q19842222)

   lcquad_train_23123: What is the state of India that contain the word "uttarakhand" in its name ?

   expected_entity: state of India (Q12443800 redirected from Q13390680)

3. DBLP: We found that the DBLP dataset contains a large amount of very ambiguous references to authors (for instance: C. Li is a very common name) and also has other string objects referring to entities. Specific to DBLP-QuAD is the absence of context for many of these cases. This makes a disambiguation clealry impossible.

   dblp_train_Q0063: What are the papers written by the person C. Li?

   expected_entity: https://dblp.org/pid/22/3837

   dblp_valid_Q0669: Does the publication 'A complete study of electroactive polymers for energy scavenging: modelling and experiments' have bibtex type Article?

   expected_entity: http://purl.org/net/nknouf/ns/bibtex#Article (this entity and the required triple of the SPARQL query is not part of the DBLP-RDF KG)

   dblp_train_Q4897: Did the authors of the publication 'Cloud of Line Distribution and Random Forest Based Text Detection from Natural/Video Scene Images' also publish a publication in ECOC?

   expected_entity: ECOC (the entity for ECOC is not contained in DBLP-RDF and the required SPARQL query references the entity only by matching to the RDF label)

4.2.5. Other special cases

This section deals with special set of cases, that does not belong to previous cases. Firstly, we exclude all questions from the dataset that do not need any entities for answering the question. With the remaining questions the entity matching operation is performed. Some of the questions in no match category were due to issues specific to the dataset creation. The first three cases were common to LC-QuAD 2.0 and SMART 2022 dataset, while the fourth case is specific to DBLP-QuAD. We discuss them below.

1. Spelling errors with entity names

   lcquad_train_7512: What award did <u>Gyorgy Ligeta</u> receive on 2000-0-0?

   expected_entity: györgy ligeti

   lcquad_train_48964: When was Pablo Picasso's partnership with <u>Fernade Oliver</u> over?

   expected_entity: fernande olivier

   lcquad_train_26637: Which is a used metre of <u>lliad</u>?

   expected_entity: iliad

2. Incomplete or empty questions

   lcquad_train_9124: None

   expected_entity: anders behring breivik

   lcquad_train_14075: What is it?

   expected_entity: sawda bint zama

   lcquad_train_13863: This sentence makes no sense.

   expected_entity: story musgrave

   smart2022_train_dbp_49907: n/a

   expected_entity: War on Terror, Mass murder

3. Entity not required for query formulation

   lcquad_train_6564: Where was <u>Augustus II the Strong</u> buried?

   expected_entity: katholische hofkirche

   lcquad_train_12811: Who is the participant on the event of 1945-9-2?

   expected_entity: günter grass

4. Abbreviated names in DBLP-QuAD dataset: author names do not follow a standard format.

   dblp_train_Q0007: What is the Wikidata ID of Zachariasen, M.?

   author_name: Martin Zachariasen

   dblp_train_Q0049: What is the webpage of the person Vidal, Vincent?

   author_name: Vincent Vidal

   dblp_train_Q0092: What are the papers written by Octaviano, M. V.?

   author_name: Manolito V. Octaviano

4.3. Discussion

So far, we looked at different cases of entity representation in natural language questions for KGQA datasets. We also found issues both common and unique to datasets and KGs. Understanding them helps build a more robust KGQA system better at handling these issues.

Firstly, we found that building a list of context labels (abbreviations, partial labels) for entities can be beneficial in significantly improving EL. It can help match variation of entity representations as demonstrated in the case of Partial Match. Table 2 shows that a huge number of entity matches were attributed to this case across all datasets. For the DBLP-RDF KG, building a dictionary of combination of author name abbreviations helped double the number of entity matches (partial match count is almost equal to exact match count). The No Match for entity contains cases (2, 3 and 4) that can be resolved to some extent through additional labels such as region, language and grammatical: parts-of-speech context. If built extensively, the context labels can also help mitigate spelling errors of entities.

Trained approach for entity detection can be another way to resolve ambiguity, by using augmented data to train and identify entities of different forms of representation. This would require a lot of training data and can work well for domain specific systems. But with KGQA datasets containing heterogeneous, open domain questions there is always a possibility of finding out-of-vocabulary entities.

KGs are updated fairly regularly to keep up with the new flow of information. While URI updates should be avoided in general, mismatches can be mitigated with constant updation of the entity dictionary (URI and context labels) to keep up with KG updates. Lastly, incorrect or incomplete input data leads to inaccurate results that cannot be fairly evaluated. Care has to be taken into preparation of these datasets and exclude any records that contain erroneous data from analysis.

5. Summary

Earlier studies conducted on EL focus on specific KGs like Wikidata [10] or on a different domain benchmark datasets (blog posts and news articles) [13]. In this paper, we provided an in-depth analysis of different entity forms and its challenges to EL in KGQA. We considered multiple benchmark datasets from three different KGs for our study. Each dataset was pre-processed to identify entity surface forms in the question and compared in its closeness to the KG representation. The findings were categorised and demonstrated with examples. The issues of concern were discussed with possible solutions to help resolve them. The goal of the study was to help the community understand the different challenges posed by surface form entities for EL. Equipped with this knowledge researchers can build more efficient and adaptable KGQA systems.

References

[1] W. Shen, J. Wang, J. Han, Entity linking with a knowledge base: Issues, techniques, and solutions, IEEE Transactions on Knowledge and Data Engineering 27 (2014) 443–460.

[2] X. Yin, D. Gromann, S. Rudolph, Neural machine translating from natural language to sparql, 2019. arXiv:1906.09302.

[3] K. Singh, A. S. Radhakrishna, A. Both, S. Shekarpour, I. Lytra, R. Usbeck, A. Vyas, A. Khikmatullaev, D. Punjani, C. Lange, et al., Why reinvent the wheel: Let's build question answering systems together, in: Proceedings of the 2018 world wide web conference, 2018, pp. 1247–1256.

[4] S. Wang, J. Jiao, Y. Li, X. Zhang, Z. F. 0002, Answering questions over rdf by neural machine translating, in: K. L. Taylor, R. Gonçalves, F. Lécué, J. Yan (Eds.), Proceedings of the ISWC 2020 Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 19th International Semantic Web Conference (ISWC 2020), Globally online, November 1-6, 2020 (UTC), volume 2721 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020, pp. 189–194. URL: http://ceur-ws.org/Vol-2721/paper549.pdf.

[5] P. Christmann, R. Saha Roy, G. Weikum, Question entity and relation linking to knowledge bases via clocq, in: 2nd Semantic Reasoning Evaluation Challenge 3rd SeMantic Answer Type, Relation and Entity Prediction Tasks Challengeand, CEUR-WS. org, 2022, pp. 33–47.

[6] A. Jiomekong, B. Tagne, Towards an approach based on knowledge graph refinement for tabular data to knowledge graph matching, Semantic Web Challenge on Tabular Data to Knowledge Graph Matching (SemTab), CEUR-WS. org (2022).

[7] A. Sakor, K. Singh, A. Patel, M.-E. Vidal, Falcon 2.0: An entity and relation linking tool over wikidata, in: Proceedings of the 29th ACM International Conference on Information & Knowledge Management, 2020, pp. 3141–3148.

[8] L. Wu, F. Petroni, M. Josifoski, S. Riedel, L. Zettlemoyer, Scalable zero-shot entity linking with dense entity retrieval, arXiv preprint arXiv:1911.03814 (2019).

[9] N. Steinmetz, Entity linking for kgqa using amr graphs, in: European Semantic Web Conference, Springer, 2023, pp. 122–138.

[10] C. Möller, J. Lehmann, R. Usbeck, Survey on english entity linking on wikidata: Datasets and approaches, Semantic Web 13 (2022) 925–966.

[11] N. Ngomo, 9th challenge on question answering over linked data (qald-9), language 7 (2018) 58–64.

[12] M. Dubey, D. Banerjee, A. Abdelkawi, 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 18, Springer, 2019, pp. 69–78.

[13] M. Van Erp, P. Mendes, H. Paulheim, F. Ilievski, J. Plu, G. Rizzo, J. Waitelonis, Evaluating entity linking: An analysis of current benchmark datasets and a roadmap for doing a better job, in: Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16), 2016, pp. 4373–4379.