

No One is Perfect: Analysing the Performance of Question Answering Components over the DBpedia Knowledge Graph

Kuldeep Singh
Fraunhofer IAIS
Sankt Augustin, Germany
kuldeep.singh@iais.fraunhofer.de

Ioanna Lytra
University of Bonn & Fraunhofer
IAIS
Bonn, Germany
lytra@cs.uni-bonn.de

Arun Sethupat Radhakrishna
University of Minnesota
Minnesota, USA
sethu021@umn.edu

Saeedeh Shekarpour
University of Dayton
Dayton, USA
sshekarpour1@udayton.edu

Maria-Esther Vidal
TIB & L3S Research Centre
Hannover, Germany
lytra@cs.uni-bonn.de

Jens Lehmann
University of Bonn & Fraunhofer
IAIS
Bonn, Germany
jens.lehmann@cs.uni-bonn.de

ABSTRACT

Question answering (QA) over knowledge graphs has gained significant momentum over the past five years due to the increasing availability of large knowledge graphs and the rising importance of question answering for user interaction. DBpedia has been the most prominently used knowledge graph in this setting and most approaches currently use a pipeline of processing steps connecting a sequence of components. In this article, we analyse and micro evaluate the behaviour of 29 available QA components for the DBpedia knowledge graph that were released by the research community since 2010. As a result, we provide a perspective on collective failure cases, suggest characteristics of QA components that prevent them from performing better and provide future challenges and research directions in the field.

1 INTRODUCTION

In the era of Big Knowledge, Question Answering (QA) systems allow for responding natural language or voice-based questions posed against various data sources, e.g. knowledge graphs, relational databases or documents [6, 17]. Particularly, with the advent of open knowledge graphs (e.g. DBpedia [1], Freebase [3] and Wikidata [41]), question answering over structured data gained momentum and researchers from different communities, e.g. semantic web, information retrieval, databases and natural language processing have extensively studied this problem over the past decade [19, 38, 43]. Thus, since 2010, more than 62 QA systems have been published, and DBpedia is the underlying knowledge graph in 38 of them [38]. Those systems usually translate natural language questions to a formal representation of a query that extracts answers from the given knowledge graph. The analysis of the architecture of these QA systems over DBpedia shows that the QA system architectures share similar tasks on the abstract level. These tasks include Named Entity Recognition and Disambiguation (NER and NED), Relation Linking (RL), Class Linking (CL), answer type identification, dependency parsing and Query Building (QB) [10].

For instance, for the question “What is the time zone of New York City?”, an ideal QA system over DBpedia generates a formal

representation of this question, that is a formal query (here expressed as SPARQL¹), which retrieves answers from the DBpedia endpoint² (i.e. `SELECT ?c {dbr:New_York_City dbo:timeZone ?c. }`). During this process, a QA system performs successive QA tasks. In the first step (i.e. NED), a QA system is expected to recognise and link the entity present in the question to its candidate(s) from DBpedia (i.e. mapping New York City to `dbr:New_York_City`³). The next step is RL, in which QA systems link the natural language predicate to the corresponding predicate in DBpedia (i.e. mapping time zone to `dbo:timeZone`⁴). Finally, the QB component obtains the linked resources and predicates from the previous steps to formulate a SPARQL query.

Research Objective. Several independent QA components for various QA tasks (e.g. NED and RL) have been released by the research community. These components are earlier reused in QA frameworks such as openQA [18], QALL-ME [12], OKBQA [16] and Frankenstein [31] to build QA systems in collaborative community efforts rather building a system from scratch. Recent empirical studies have revealed that albeit overall effective, the performance of monolithic QA systems and QA components depends heavily on the features of input questions [28, 33], and not even the combination of the best performing QA systems or individual QA components retrieves complete and correct answers [6]. In order to advance the state of the art in building QA systems and explore future research directions, it is important to get insights into the strengths and weaknesses of the range of existing QA components that can be reused in building QA systems. Therefore, the goal of this evaluation study is to analyse the performance of existing reusable QA components implementing the aforementioned QA tasks with respect to different question features, and identify, thus, common behaviours, collective failures and directions for future QA research.

Approach. In this paper, we aim at putting main properties of existing QA approaches over DBpedia in perspective and provide strong evidence and conclusive insights about the parameters that affect the performance of state-of-the-art QA approaches. We have collected 29 QA components that implement various QA tasks, and evaluate each component over 3,000 questions from the LC-QuAD dataset [35]. 19 out of 29 QA components are accompanied by peer reviewed publications, and other 10

¹<https://www.w3.org/TR/rdf-sparql-query/>

²<http://dbpedia.org/sparql>

³Prefix dbr is bound to <http://dbpedia.org/resource/>

⁴Prefix dbo is bound to <http://dbpedia.org/ontology/>

		TagMe	DBpedia Spotlight		TagMe	DBpedia Spotlight	
Upper/Lower Case	Q1: When was university of edinburgh founded?	✓	✗	Number of Words in Entities	Q1: List all the chairmen of Vancouver Whitecaps FC ?	✗	✓
	Q2: When was University Of Edinburgh founded? dbr:University_of_Edinburgh	✓	✓		dbr:Vancouver_Whitecaps_FC		
Explicit/Implicit Entity	Q1: Who is the current prime minister of Australia ?	✓	✓		Q2: Who has been a manager of english under twenty one football team ?	✗	✗
	Q2: Who is the current Australian prime minister? dbr:Australia	✗	✗		dbr:England_national_under-21_football_team		
Abbreviations	Q1: What is the time zone of New York City ?	✓	✓				
	Q2: What is the time zone of NYC ? dbr:New_York_City	✓	✗				

Figure 1: Performance of Top-2 NED Components on Specific Questions. TagMe and DBpedia Spotlight are the top-2 NED components in average on the LC-QuAD dataset (as shown in Table 1). However, their behaviour varies with regard to question features, for which entities in lowercase vs. uppercase, including abbreviations, having implicit vs. explicit mappings and with varying number of words have to be mapped to the DBpedia knowledge graph entities.

components are openly available to be reused for building QA systems. These components are evaluated using the questions range from simple to complex and vary in the number of entities and relations as well as expressiveness. Overall, 59 variables or question features are observed during the evaluation of these questions in the 29 QA components. The outcome of this evaluation uncovers characteristics of the studied QA components that allow for the explanation of their diverse behaviour. First, there are certain questions that none of the components or their combinations are able to answer, e.g. all NED components fail to answer some questions with different capitalisation. Second, we also have observed that best performing QA components for a particular QA task are not the best for all the question features. For example, TagMe [11] is the overall best component evaluated by Singh et al. [33] for the NED task, but it is not the best component for all the question features reported in this paper.

Contributions. Our work contributes with four folds: First, results and analysis from a micro evaluation of 29 QA components based on 59 question features using over 3,000 questions per component are presented. Second, we provide a perspective on collective failure cases of state-of-the-art QA components and pitfalls of these components. Third, we collect insights suggesting the characteristics of the QA components that prevent them from correctly and completely performing their corresponding QA task. Finally, challenges and research directions required to be addressed in order to advance the state of the art are discussed. The remainder of the paper is structured as follows. We illustrate the diverse behaviour exhibited by existing QA components in Section 2. Then, Section 3 discusses related work and analyses drawbacks of existing QA benchmarkings and evaluations. In Section 4, the experimental configuration is detailed, and the results of the empirical evaluation are reported in Section 5. We analyse the collective failures of QA components in Section 6. In Section 7, we discuss and provide insights about the observed results, and Section 8 concludes with an outlook on future work.

2 MOTIVATING EXAMPLES

One of the major tasks in a QA system is Named Entity Disambiguation (NED). Although there is a high number of implementations for the NED task [11, 14, 20], none of the implementations consistently performs best on a variety of input

questions. Various aspects of the input questions, e.g. length or complexity, might influence the performance of NED components. Our motivating example demonstrates the dependency of the performance of two well-known NED implementations for various expressions of natural language input questions (cf. Figure 1). For example, with respect to the question “When was Edinburgh University founded?”, an ideal NED component targeting DBpedia⁵ [1] as underlying knowledge graph will identify the segment Edinburgh University as a named entity and afterwards link it to suitable candidate(s) in the knowledge graph, e.g. `dbr:University_of_Edinburgh`. TagMe⁶ [11] and DBpedia Spotlight⁷ [20] are NED components evaluated in [33] as the top 2 (out of 18) components, assessed using the LC-QuAD [35] dataset. The precision, recall and F-score for these components are reported in Table 1. In the following paragraphs, we show the dependency of their performance on the input question.

Table 1: Top-2 NED Components on DBpedia. The performance of TagMe and DBpedia Spotlight is reported based on the LC-QuAD dataset. They are the best performing components across 3,253 questions; however, the performance of these NED components varies depending on the question features (cf. Figure 1).

QA Component	QA Task	Precision	Recall	F-score
TagMe	NED	0.69	0.66	0.67
DBpedia Spotlight	NED	0.58	0.59	0.59

Impact of Uppercase/Lowercase Entities. For the question “When was Edinburgh University founded?”, TagMe and DBpedia Spotlight are able to successfully recognise and disambiguate the entity `dbr:University_of_Edinburgh`. If the given question is rephrased to “When was university of edinburgh founded?”, in which the entity mention is written in *lowercase letters*, TagMe continues to identify and disambiguate the correct entity whereas

⁵<https://wiki.dbpedia.org/>

⁶<https://services.d4science.org/web/tagme/demo>

⁷<https://www.dbpedia-spotlight.org/demo/>

DBpedia Spotlight fails to do so. However, if the question is further rephrased to “When was University Of Edinburgh founded?”, in which the entity segment appears in *uppercase letters*, both NED components successfully disambiguate the entity.

Impact of Implicit/Explicit Entities. Let us consider another question, namely, “Who is the current prime minister of Australia?”. In this question, the correct DBpedia mention for the named entity Australia is `dbr:Australia`. TagMe and DBpedia Spotlight can correctly disambiguate this entity. However, if the question is rephrased to “Who is the current Australian prime minister?”, both components fail to link the entity Australian to `dbr:Australia`. This is due to the fact that the entity is no longer an exact match of its DBpedia label `dbr:Australia` (i.e. vocabulary mismatch problem [29]), and therefore, these components need to rewrite or reformulate the input question.

Impact of Abbreviations in Entities. The abbreviations presented in the question have an impact on the performance of the NED components. For the question “What is the time zone of New York City?”, both TagMe and DBpedia Spotlight can disambiguate the entity New York City to `dbr:New_York_City`. However, when this question is rephrased as “What is the time zone of NYC?”, TagMe succeeds while DBpedia Spotlight fails.

Impact of Number of Words in Entities. For the fourth question feature, regarding the question “List all the chairmen of Vancouver Whitecaps FC?”, the segment Vancouver Whitecaps FC is expected to be linked to `dbr:Vancouver_Whitecaps_FC`. Although the overall performance of DBpedia Spotlight is lower than TagMe, it succeeds in linking the entity while TagMe fails to recognise and disambiguate Vancouver Whitecaps FC to its target entity. In this question, the entity Vancouver Whitecaps FC has three words. We observe that if the number of words in the entity increases further, then DBpedia Spotlight also fails to identify entities. For instance, for the question “Who has been a manager of english under twenty one football team?”, both components fail to disambiguate the entity english under twenty one football team (6 words) to its target DBpedia entity.

These observations lead us to assume that certain characteristics of input questions such as length, variety of expressions and number of words in entities impact the performance of QA components. In other words, a fine-grained analysis of performance indicates that the overall performance of a component is not representative when expressed as an average performance value on all the questions of a dataset because we observe a correlation between the input question characteristics and the success rate.

3 RELATED WORK

QA Systems and Evaluation. Question answering systems for factoid questions over different publicly available Knowledge Graphs (KGs) have been evaluated using various datasets/benchmarks. SimpleQuestion [4] and WebQuestions [2] are popular datasets for evaluating systems that use Freebase as underlying KG. Most of the state-of-the-art QA systems over Freebase have been evaluated using the SimpleQuestions and WebQuestions dataset [2, 7, 42]. For DBpedia, Question Answering over Linked Data (QALD) benchmark series provides widely used datasets for evaluating QA systems [36]. The number of questions in the QALD datasets ranges from 50 to 350 across its different editions. Therefore, state-of-the-art QA systems over DBpedia have been evaluated mostly on limited number of questions (50–100) [6, 38, 43]. LC-QuAD [35] is another recently released dataset

that contains 5,000 questions for DBpedia and has been used for evaluating question answering systems [9] as well.

Although the aforementioned benchmarks have been extensively used to evaluate QA systems as well as QA related components, no focus has been set on analysing the results with respect to the different types of questions included in these benchmarks. Therefore, across datasets and KGs, QA systems and components have reported their overall accuracy *on average over all the questions present in the dataset* (i.e. by performing macro evaluation), which does not provide many insights about the weaknesses and strengths of a particular QA component and system when considering different types of questions.

Hybrid Question Answering System. Cui et al. [6] introduced the concept of hybrid question answering systems in 2017. The authors combine their proposed KBQA QA system with existing monolithic QA systems to build hybrid QA systems. The evaluation shows that this increases the overall accuracy of the hybrid QA system for the 99 questions from QALD-3 [5] dataset. Singh et al. [33] extend the applicability of hybrid QA systems to component level and combine different components per question answering task to build QA systems. The work has been evaluated using the 204 questions of QALD-5 [36] and more than 3,000 questions of the LC-QuAD [35] dataset. In these studies, the performance of hybrid QA systems and components has increased, however, the characteristics of the unanswered questions have not been further analysed. It is also not clear for which types of questions hybrid QA systems and components find limitations. Therefore, in order to build hybrid QA systems collaboratively, it is also important to understand the strengths and weaknesses of different components that are used to build hybrid QA systems.

4 EVALUATION SETTINGS

In this section, we describe the evaluation settings of the experiments reported in this paper. In Section 4.1, we list the QA components considered in our experiments. Then, Section 4.2 describes the experimental setup and employed metrics for assessment. Section 4.3 explains the benchmarking procedure. Finally, Section 4.4 presents all features derived from the questions.

4.1 Components for Evaluation

The *Frankenstein framework* [33] was employed as the underlying platform to run and evaluate QA components for our experiments since compared to the state of the art (i.e. openQA [18], QALLME [12] and OKBQA [16]) it includes a higher number of QA components. Frankenstein wraps the functionality of public APIs of these components and integrates them in its core QA pipeline as microservices. Therefore, it does not have any impact on the functionality of these components. However, most of the QA components when reused in QA frameworks target high precision rather than high recall per question. For instance, in the case of TagMe⁸ and for the given question “Who is the wife of Barack Obama?” three entities with different confidence scores are returned (Barack Obama, First_Lady_of_the_United_States and World_Health_Organization). Frankenstein picks entities with confidence score more than 0.50 which leads to the selection of only one entity i.e. Barack Obama (correct in case of this question to formulate SPARQL query), thus maximising the precision. Frankenstein currently integrates 29 QA components including 11 NER, 10 NED, five RL and two QB components that have

⁸<https://services.d4science.org/web/tagme/demo>

been used to build QA systems⁹. Except for 11 NER, rest of the NED components implicitly perform the NER task and provide directly the disambiguated URLs which are required to formulate the SPARQL query for an input question. Therefore, we focus on the NED task excluding the independent evaluation of NER components. However, all 11 NER components are used jointly with the AGDISTIS [37] disambiguation component to evaluate how well these NER components recognise entities that can act as input for a disambiguation component such as AGDISTIS.

Besides the most common QA tasks such as NED, RL and QB, in specific questions, it is necessary to map to the ontology classes¹⁰ present in the question. For the exemplary question “Which comic characters are painted by Bill Finger?”, its intermediate representation using DBpedia (“Which dbo:ComicsCharacter are dbo:creator by dbr:Bill_Finger?”) requires recognising and linking to the expected DBpedia class dbo:ComicsCharacter as one of the inputs for constructing the corresponding SPARQL query. Frankenstein includes two Class Linking (CL) components. To the best of our knowledge, there are no other independent components (for DBpedia) except the 29 components integrated in Frankenstein for NED (20), RL (5), QB (2) and CL (2)¹¹.

4.2 Evaluation Setup

We executed our experiments on ten virtual servers, each with eight cores, 32 GB RAM running on the Ubuntu 16.04.3 operating system. We have reused the open source implementation of Frankenstein¹² for executing the different QA components and Stardog v4.1.3¹³ as RDF (Resource Description Framework)¹⁴ datastore to store all data generated in the experiments.

Knowledge Graph. DBpedia¹⁵ [1] was our underlying background knowledge graph. It has 5.6 million entities and 111 million (subject,predicate,object) triples. It requires 14.2 GB storage.

Test Data. QALD¹⁶ and LC-QuAD [35] are two datasets of *factoid* questions available to benchmark QA systems and components over DBpedia. The total number of questions in the QALD series is relatively low (350 questions at most in the 5th edition) and the questions do not have much diversity. Compared to QALD, the newly released LC-QuAD has 5,000 questions, the highest number of benchmarking questions available for DBpedia [35]. The fully annotated LC-QuAD dataset (with entity, relation and class labels) has over 5,000 entities and more than 615 predicates covered in its questions. The questions are also diverse in expressiveness, as 20% of questions are simple (questions with only one entity and one relation) and 80% are complex questions. However, only 3,253 questions from this benchmark return answers from the latest DBpedia version¹⁷ because the dataset was created using a previous DBpedia version 16-04. Since all components integrated in Frankenstein use the latest version of DBpedia for their functionalities, we have considered only this subset of LC-QuAD. These 3,253 questions were evaluated against the 29 Frankenstein components independently, resulting into 94,337 total questions executed using the Frankenstein framework. To the

best of our knowledge, this is the first study over DBpedia with a high number of questions and components. Frankenstein platform generates a KG of results using RDF (Resource Description Framework)¹⁸ format [31] for further analysis. The execution of all questions over all available components generates approximately 4.9GB RDF data. The experimental data, the questions, source code for experiments and detailed results can be found in our open source repository for reusability and reproducibility¹⁹.

Metrics: The following evaluation metrics have been employed: *i)* Micro Precision (MP): For a given component, the ratio of correct answers vs. total number of answers retrieved for a particular question. *ii)* Precision: The average of the Micro Precision over all questions by a component. *iii)* Micro Recall (MR): The number of correct answers retrieved by a component vs. gold standard answers for the given question. *iv)* Recall (R): The average of Micro Recall over all questions for a given component. *v)* Micro F-score (MF): Harmonic mean of MP and MR for each question. *vi)* F-score: Harmonic mean of P and R for each component. *vii)* Processed Question: Question for which a component has non zero F-score. *viii)* Unprocessed Question: Question for which a component has zero Micro F-score. *ix)* Baseline Value: The highest F-score value for a particular question feature. *x)* Baseline Component: The component with the highest F-score value for a particular question feature. *xi)* State-of-the-art F-score: The highest F-score value per task on average over all questions.

4.3 Component Benchmarking

In a particular QA task, the performance of each component is measured initially by calculating the Micro F-score per question, and then computing the Macro F-score representing the overall performance. For instance, consider the question “Which comic characters are painted by Bill Finger?” from the LC-QuAD dataset. The corresponding SPARQL query for this question is:

```
SELECT DISTINCT ?uri WHERE
{?uri <http://dbpedia.org/ontology/creator>
<http://dbpedia.org/resource/Bill_Finger> . ?uri
<https://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://dbpedia.org/ontology/ComicsCharacter>}
```

In order to benchmark the NED components, the output URIs of the disambiguated entities produced by each component must be compared against the correct DBpedia entity (dbr:Bill_Finger in this case). Afterwards, for each component Micro Precision, Micro Recall and Micro F-score values are calculated. A similar procedure has been followed for the RL and CL components, whose outputs are compared to the correct URIs (i.e. dbo:creator for the RL and dbo:ComicsCharacter for the CL task). For each of the QB components, given a correct set of URIs as input, the generated SPARQL query is compared to the benchmark SPARQL query for the given question by comparing the answers the SPARQL queries retrieve from DBpedia. A similar component benchmarking procedure has been followed in [23, 32].

4.4 Question Features

Question classification based on features (e.g. question length, POS tags, head words etc.) has been a continuous field of research after the first edition of TREC challenge [40] for open domain QA. However these features are explicitly used for answer type classification [15] and not for benchmarking/evaluating QA systems.

⁹Full list of components: <https://github.com/WDAqua/Frankenstein/blob/master/Component%20List.csv>

¹⁰<https://www.w3.org/2002/07/owl#Class>

¹¹Based on the research published till 31/07/2018

¹²<https://github.com/WDAqua/Frankenstein>

¹³<https://www.stardog.com/>

¹⁴<https://www.w3.org/RDF/>

¹⁵DBpedia version 2016-10

¹⁶<https://qald.sebastianwalter.org/>

¹⁷<http://dbpedia.org/sparql> (version 2016-10)

¹⁸<https://www.w3.org/RDF/>

¹⁹HiddenforBlindReview

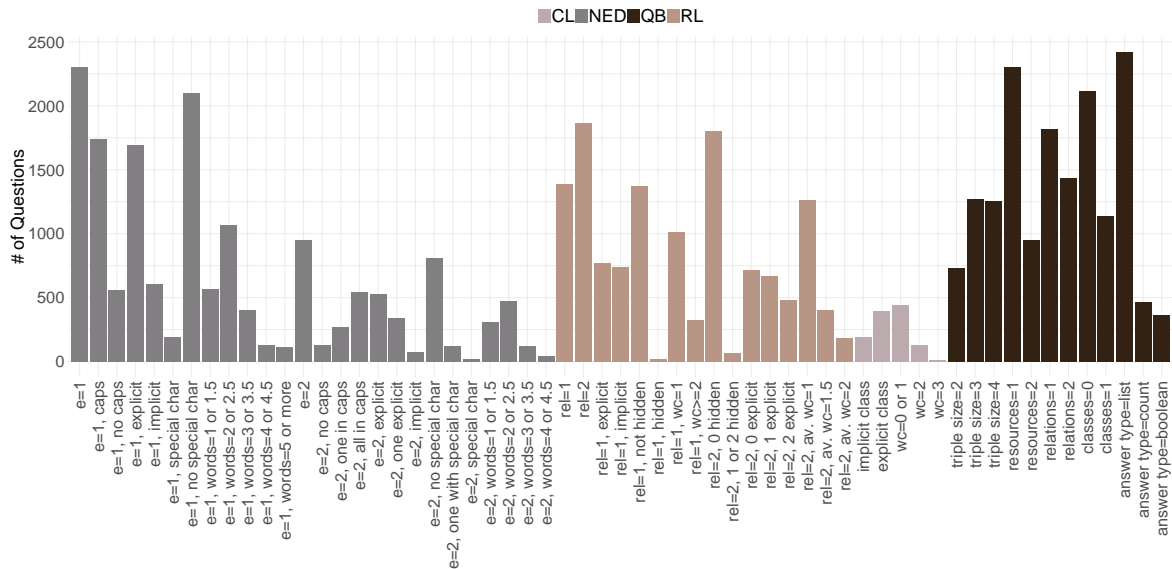


Figure 2: Distribution of Questions per Question Feature. For evaluating components fulfilling different QA tasks, various question features have been considered. 59 features have been extracted in total – 26 for NED, 16 for RL, 5 for CL and 12 for QB. Some features are considered in several questions, and some questions are characterised by several features.

In QA over KGs, NED, RL, CL and QB are concerned with the formation of final SPARQL queries and not with finding the right source of the answer. Therefore, for analysing the performance of QA components, and in order to provide a more fine-grained evaluation, we categorise questions per feature and evaluate the single QA components for these question features (cf. Figure 2 for the distribution of questions per question feature). All these question features have been automatically extracted.

4.4.1 Question Features for NED Task. For the NED task, we consider the following features of input questions.

Number of Entities. For a given question, we consider the number of entities included in the question. For example, in the question “Which comic characters are painted by Bill Finger?”, there is one entity (Bill Finger), and therefore this question is classified as “question with entity=1”.

Number of Words in Entities. The number of words of an entity is the next feature. For the exemplary question “Which comic characters are painted by Bill Finger?”, the number of words in the entity label (i.e. Bill Finger) is two; hence, this question is annotated with the feature “entity=1, number of words in entities=2”. For questions having several entities, the average number of words in all labels of entities is considered.

Implicit/Explicit Entities. The entity of a question is explicit if there is an exact match between the segment and the DBpedia resource URI label (meaning there is no vocabulary mismatch). For the exemplary question, the segment Bill Finger matches the label of `dbr:Bill_Finger`, while for the question “What religions do politicians in the Korean Minjoo Party follow?”, the segment Korean Minjoo Party and the desire resource `dbr:Minjoo_Party_of_Korea` have a slight vocabulary mismatch. We consider the explicit/implicit entities as another feature. Please note that questions with two entities might have three different statuses 1) both entities are explicit, 2) at least one of them is explicit and 3) both are implicit.

Case Sensitivity of Entities. The use of uppercase or lowercase for entities is another question feature. If all words in the segment

associated to an entity of a given question are uppercase, then we annotate the question as “uppercase”, otherwise as “lowercase”. For questions containing two entities, three different annotations for this feature are: 1) both entities are in uppercase, 2) at least one entity is in uppercase and 3) both entities are in lowercase.

Special Characters in Entities. If the entity segment contains ASCII punctuation, symbols or numbers, the question is annotated as a question with special character(s). For example, the question “What is the route junction of Rhode Island Route 15?” contains a number in the entity segment Rhode Island Route 15, therefore it is considered as a question with special characters. For questions with two entities, three annotations are possible: 1) both entities have special characters, 2) at least one has special character and 3) none has special characters.

4.4.2 Question Features for RL Task. The features for the RL task are similar to the NED task except that we excluded (i) case sensitivity and (ii) inclusion of special characters, because there were no relations with such features in the LC-QuAD questions.

Number of Relations. The number of relations in a given question is the first question feature.

Number of Words in Relation. The number of words in the relation segment is the next feature. For questions with two relations or more, this feature is computed as the average number of words in the relation segments.

Explicit/Implicit Relation. Similar to the NED task, this feature refers to whether the relation mention is explicit or implicit. In the exemplary question “Which comic characters are painted by Bill Finger?”, the relation segment painted by does not explicitly match to the DBpedia relation `dbo:creator`, therefore, this question is considered as implicit relation question.

Hidden Relation. Consider the question “Which companies have launched a rocket from Cape Canaveral Air Force station?” which contains two relations (i.e. `dbo:launchSite` and `dbo:manufacturer`). For `dbo:launchSite`, there is a mention in the question (launched), while, there is no mention for the relation `dbo:manufacturer`. We characterise questions without

a natural language segment as “questions with hidden relations”. For questions with two relations, possible annotations are: 1) one relation is hidden and 2) at least one is hidden.

4.4.3 Question Features for CL Task. All LC-QuAD questions contain at most one class. For the questions with a single class, we considered the following features: 1) explicit/implicit mention of class and 2) number of words in the class segment.

4.4.4 Question Features for QB Task. For the evaluation of the QB components, we take into consideration the following features:

Number of Triple Patterns. This feature refers to the number of basic triple patterns in the SPARQL query corresponding to the question.

Number of Relations. It means the number of relations in the SPARQL query.

Number of Entities. It relates to the number of entities in the SPARQL query.

Number of Classes. This feature refers to the number of classes in the SPARQL query.

Answer Type. There are three types of possible types: list, number or boolean. For instance, the question “Which comic characters are painted by Bill Finger” has a list as expected answer.

5 EVALUATION RESULTS

We pursue the following research questions in our experiments. **RQ1)** What is the performance of QA components depending on the question features? **RQ2)** Which QA components exhibit similar behaviour across the different questions and question features? **RQ3)** What is the impact of combining QA components in the overall performance? This section is followed by the results of the empirical experiments conducted to address our research questions. Section 5.1 represents the experiments with respect to RQ1. In Section 5.2 and Section 5.3 we discuss our observations and answers reached for RQ2 and RQ3 respectively.

5.1 Evaluating QA Components

In this experiment, we evaluate 29 components for various QA tasks based on the question features described in Section 4.4. We then report the results for each QA task independently.

5.1.1 Performance of NED Components. The heatmap in Figure 3 shows the performance (i.e. F-score) for 20 NED components across questions with different features. For example, TagMe [11] which in previous literature reported F-score value 0.67 [33], has fluctuating F-scores here.

Performance varies across question features. The baseline F-score value varies across different question features. For example, for the question feature “number of entities=1, and contains special character” (e.g. question “What is the route junction of Rhode Island Route 15?” where the entity label Rhode Island Route 15 contains special characters), TagMe [11] is the baseline component with F-score value of 0.72. However, for questions such as “What religions do politicians in the Korean Minjoo Party follow?”, in which the number of entities is one and the entity is implicit (the segment Korean Minjoo Party is mapped to `dbr:Minjoo_Party_of_Korea`), AmbiverseNED 2.0²⁰ [14] is the baseline component, but with lower baseline value – 0.41.

Impact of Number of Entities. When the number of entities in the question is two (e=2), 16 NED components illustrated an




Figure 3: Comparison of F-scores per Question Feature for QA Components Performing the NED Task. Darker colour indicates higher F-score. While TagMe and AmbiverseNED 2.0 demonstrate high performance, few components such as Babelfy, Ambiverse 1.0 and Meaning-Cloud deliver poor results for the LC-QuAD questions. These results suggest that the behaviour of best performing NED components is impacted by the question features.

increase in the F-score values compared to questions with one entity. For such questions, AmbiverseNED 2.0 is the baseline component with F-score 0.79. When the number of entities is one, all components have F-scores less than 0.67 which is the state-of-the-art F-score for the NED task. In this case, TagMe sets the baseline F-score to 0.65.

Impact of Number of Words in Entities. For questions which contain one entity, if the number of words (“wc” in heatmap) in the entity increases from one to two, 19 out of the 20 NED components demonstrate an increase in the F-score values. However, when the number of words in the entity further increases to three, 16 components demonstrate decrease in F-scores. For four or more words in the entity, only OntotextNED²¹ shows no variation in the F-score compared to questions with entities containing three words; for all other components, F-score is further decreased. In questions containing two entities, the highest F-score value remains above 0.75 independently of the average number of words in the entities.

Impact of Special Characters in Entities. Except for TagMe and DandelionNED²², for all other NED components we observed a decrease in the F-score values in the case that the spotted entities contain special characters.

Impact of Character Cases in Entities. The character cases significantly impact the performance of the NED components. The majority of QA components demonstrate a steep decrease in the F-score values when the entities appear in lowercase, independent of the number of entities in the question. For example, in

²⁰<https://www.ambiverse.com/>

²¹<https://ontotext.com/technology-solutions/semantic-tagging/>

²²<https://dandelion.eu/docs/api/datatxt/nex/getting-started/>

DBpedia Spotlight NED [20] (abbreviated as DBpedia NED in Figure 3) the F-score decreases from 0.64 to 0.36 when the entity character case is switched to lowercase. Similar behaviour is observed in other components as well such as AylieNER+AGDISTIS²³ [37] and StanfordNER+AGDISTIS [13, 37], for which the F-score values sharply drop from 0.50 and 0.55 to 0.15 and 0.09 respectively.

Impact of Implicit/Explicit Entities. The existence of implicit entities significantly impacts the performance of all NED components. For example, for questions with one explicit entity, the baseline F-score value is 0.78 (TagMe), whereas for implicit entities, the baseline is 0.41 (AmbiverseNED 2.0) and for the other 19 NED components is less than 0.30.

5.1.2 Performance of RL Components. The heatmap in Figure 4 shows the performance (i.e. F-score) for five RL components across questions with different features. The state-of-the-art F-score value for the RL task on LC-QuAD dataset is 0.23, achieved by the RNLIWOD component [33]. Similar to the results reported for the NED task, there is no overall baseline component and value for all question features, as illustrated in Figure 4. This experiment led us to the following observations:

Impact of Number of Relations. When the number of relations (rel=1 or 2 in Figure 4) in the question is elevated to two, the baseline F-score drops from 0.28 (ReMatch [23]) to 0.24 (RNLIWOD²⁴ component). Other components also demonstrate a decrease in performance in this case.

Impact of Number of Words in Relations. For questions containing one or two relations, an increase in the number of words in the relation label has increased the overall performance of the components.

Impact of Explicit/Implicit Relations. For questions containing one relation, for which the label is explicit, ReMatch [23] provides the baseline F-score 0.47. However, when the relation is implicit, the baseline F-score value drops to 0.16 (for RNLIWOD). ReMatch [16] reports the lowest F-score for implicit relations with F-score 0.06. In questions containing two relations, the number of implicit relations negatively impacts the performance. For instance, for questions being similar to “How many currencies are in use in places where people speak French?” which has two implicit relations (i.e. currencies and speak) that should be




Figure 4: Comparison of F-scores per Question Feature for QA Components Performing the RL Task. Darker colour indicates higher F-score. We observe that ReMatch and RNLIWOD deliver the best F-scores but reach zero F-score for some question features. These results provide evidence that the performance of RL components is affected by most of the question features.

²³<http://docs.aylien.com/docs/introduction> (AylieNER)

²⁴The component is similar to the Relation Linker of <https://github.com/dice-group/NLIWOD>

mapped to `dbo:currency` and `officialLanguage` respectively, the baseline F-score value is 0.17 (for RNLIWOD).

Impact of Hidden Relations. For questions containing a natural language label for the relation, ReMatch is the baseline with F-score 0.28. However, for the exemplary question “How many shows does HBO have?” There is no corresponding segment in the annotated dataset to be mapped to `dbo:company`. For all questions with one hidden relation, the baseline F-score is drops to 0.15. A similar drop in performance is observed in questions with two relations as well.




Figure 5: Comparison of F-scores per Question Feature for QA Components Performing the CL Task. Higher F-score is indicated with dark colour. OKBQA DM Class Linker performs better than CLSNLIWOD. However, all question features impact the performance of both components.

5.1.3 Performance of CL Components. The heatmap in Figure 5 shows the fine-grained results of our experiment for the CL task. For this, we took into account two state-of-the-art components integrated in Frankenstein. We concluded that the following question features impact their overall performance:

Impact of Explicit/Implicit Class. Consider the question “What university campuses are situated in Indiana?” for which the class label `university` is expected to be mapped to `dbo:University` through an exact match. In such cases, OKBQA DM CLS²⁵ has the baseline F-score 0.72 compared to the state-of-the-art F-score 0.52 for the CL task [33]. However, when class mentions are implicit, the baseline F-score drops to 0.45, but OKBQA DM CLS is still the best component.

Impact of Number of Words in Class. The increasing number of words in the class label has a positive impact on the performance of the components. OKBQA DM CLS remains the baseline for all question features. However, the baseline value increases with the number of words in the class label. When the class label contains one word the baseline F-score is 0.61; for two and three words, the baseline F-score increases to 0.72 and 0.94 respectively.

5.1.4 Performance of QB Components. The two QB components available in Frankenstein were evaluated based on different question features. The details of performance behaviour are illustrated in Figure 6. The state-of-the-art F-score for the QB task is 0.48 for NLIWOD QB²⁶. We noticed the following observations based on the impact of the question features on the performance:

Impact of Number of Triple Patterns. With the increase in the number of triple patterns, the performance of the QB components is downgraded in general. For queries with two triple patterns, SINA [30] is the baseline component with F-score value 0.80. When the number of triple patterns is three, SINA’s F-score drops to 0.18, and NLIWOD QB is the new baseline component with F-score 0.52. In the case of questions mapped to SPARQL

²⁵Component is similar to the CL of <http://repository.okbqa.org/components/7>

²⁶Component is based on <https://github.com/dice-group/NLIWOD>.

queries with four triple patterns, NLIWOD QB remains the baseline with F-score 0.38, while the F-score for SINA drops to 0.

Impact of Number of Resources. The increase in the number of resources (i.e. entities) in the question also inversely affects the performance of the QB components. In case of single resources, NLIWOD QB has F-score 0.54, while this value drops to 0.38 when the number of resources in the question increases to two. The F-score for SINA drops to 0 for questions with two resources.

Impact of Number of Relations. NLIWOD QB and SINA have shown dissimilar behaviour in the performance when the number of relations in the question increases. For an increase in the number of relations from one to two, the F-score value for NLIWOD QB has been elevated from 0.48 to 0.50, whereas the F-score value for SINA has dropped from 0.40 to 0.04.

Impact of Number of Classes. When there is no class in the question, NLIWOD QB is the baseline component (F-score is 0.46). For questions with one class, the baseline F-score value increases to 0.54 for NLIWOD QB which outperforms SINA in both cases.

Impact of Answer Type. NLIWOD QB remains the best component for any answer type (i.e. list, boolean and number.) For instance, the question “Which comic characters painted by Bill Finger?” expects a list of DBpedia resources as an answer. For such questions, NLIWOD QB is the baseline component with F-score 0.56. When the expected answer type is boolean, the baseline F-score value drops to 0.46, and it further drops to 0.14 for questions expecting a number as an answer. SINA cannot answer any questions with answer types number or boolean.

5.2 Clustering of QA Components

In this experiment, we pursue the second research question, that is where or when QA components exhibit similar behaviour. To study the similarity of QA components behaviour we applied clustering techniques on the evaluation results of these components based on the question features as well as overall question results. In particular, we have applied k-means clustering²⁷ To cluster QA components based on (a) performance with respect to question features (macro level) and (b) performance with respect to processed LC-QuAD questions (micro level). Figure 7a and Figure 7b visualise our clustering results. Since the number of components per QA task under evaluation differs, we were able to study this research question only for the task for which there is an adequate number of components, i.e. the NED task. In the first





Figure 6: Comparison of F-scores per Question Feature for QA Components Performing the QB Task. Higher F-score is indicated with dark colour. SINA performs better than NLIWOD Query Builder only when the two triples are included in the input question. For most of the question features F-scores remain lower than 0.6.

²⁷The optimal k was calculated using the elbow and silhouette method [27].


clustering (cf. Figure 7a), the red cluster (right) represents components with higher performance, while the black (middle) and green (left) clusters group medium and low accuracy components respectively. The “distance” between the single components can also be observed by comparing the colour scales in Figure 3. In the second clustering (cf. Figure 7b), components that appear in the same cluster (or are “closer” in the cluster) are expected to provide correct answers to similar subsets of the LC-QuAD questions. By comparing these two figures, we observe that NED components that are similar with respect to their performance over different question features are not necessarily able to answer the same questions as well (although a tendency is in general observed). For instance, TagMe NED and Babely NED [22] which appear close in the clustering of Figure 7b demonstrate different performance in Figure 7a. In addition, in order to find out which question features are more significant for clustering NED components, we have performed Principal Component Analysis (PCA). From this, we concluded that “e=2, explicit”, “e=2, words=2 or 2.5” and “e=2, no special characters” constitute the most important features. This means, in principle, that the NED components under study demonstrate very different F-scores for these features. This type of analysis allows us to group components with respect to the type of questions they are able to address successfully.

5.3 Hybrid Composition of QA Components

In this section, we address the third research question, in particular, we investigate what is the impact of combining components in the total performance. The idea of integrating QA systems was first introduced by Cui et al. [6], who combined two QA systems and observed a slight increase in the number of answered questions for 99 questions of the QALD-3 dataset. First, the input question is sent to the KBQA QA system (introduced by the authors) and if KBQA cannot answer the question this question is passed to other monolithic QA systems; the combination is known as hybrid QA system. In this experiment, we extend the concept of hybrid QA systems by *integrating at the component level* in order to study the impact of combining QA components on the total number of processed questions. We also use a bigger set of questions (3,253) from LC-QuAD compared to 99 questions from QALD used in [6]. For our experiment, we send a question to the component with highest F-score on LC-QuAD (average on all 3,253 questions), and if the question is not answered, we forward the question to the component with the next highest F-score. We continue this process for all QA tasks until the component with the lowest F-score. We calculate the accumulative number of processed questions for top-1, top-2, etc. components which are taken into account. Figure 8 illustrates the impact of combining QA components on the number of processed questions for all tasks. It can be observed that for the RL, QB and CL tasks this number continuously increases when more components are combined. However, for the NED task (where a bigger number of components is considered), the number of answered questions initially increases but it slowly saturates when more than four components are considered. This shows that combining several QA components performing the same task in a hybrid QA system does not necessarily improve the overall performance (not all the components are complementary). In addition, this observation provides an indication that most of the QA components fail for the same questions, i.e. there are several questions that cannot be addressed by any of the available components.



(a) Clustering based on question features (macro level)



(b) Clustering based on question results (micro level)

Figure 7: Clustering of NED Components. K-means clusters NED components at macro and micro level. In (a) the red cluster contains the best performing components while the black and the green clusters represent NED components with medium and low performance respectively, for questions with different features. While the components with similar performance are clustered together in (a), in (b) we observe that some are dissimilar with regard to the questions they are able to answer.

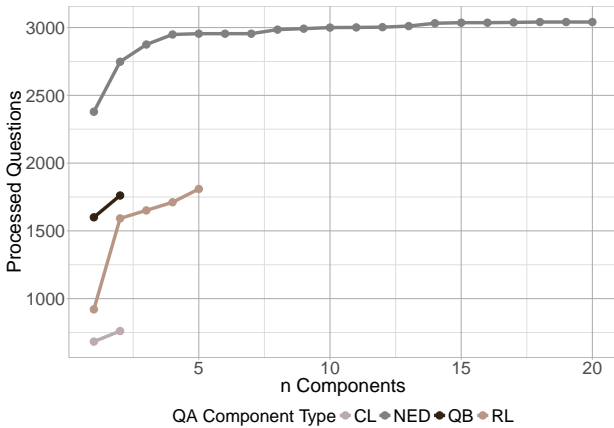


Figure 8: Total Processed Questions for the Combination of Top-n QA Components per QA Task. Combinations of two or more QA components improve the number of processed questions. Particularly, combinations of the top-4 NED are able to process a bit less than 3,000 questions; however, adding more NED components does not improve the performance significantly.

6 FAILURE ANALYSIS

There are several questions for which all components of the respective task failed to process (cf. Figure 9). We have analysed the question features for those questions per QA task in order to understand in depth the collective failure of the components.

Unprocessed Questions for NED Components. For the NED task, 210 questions cannot be processed by any of the 20 components, i.e. for which the F-score for all components is 0. As illustrated in Figure 10, 200 unprocessed questions have a single entity and just 10 questions have two entities. Out of the 200 single entity questions, 165 contain implicit entities. Furthermore, 121 questions with single entity have entity label written in uppercase letters and 79 with lowercase letters. While analysing

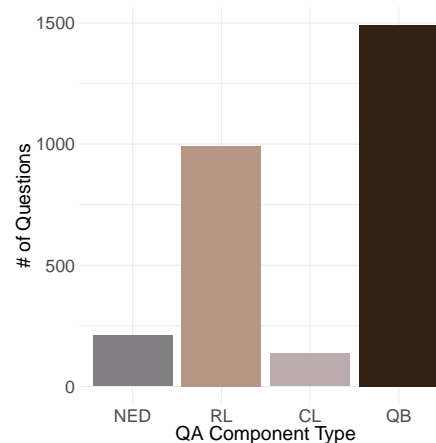


Figure 9: Number of Unprocessed Questions per QA Task. While combination of all NED and CL components are able to process the high majority of the LC-QuAD questions, the combinations of QB and RL components fails in half and one third of the questions, respectively.

these questions further, we have observed that out of 121 questions which have uppercase letters, 100 include implicit entities as well. A similar observation is valid for the 79 questions having an entity label in lowercase, out of which 64 contain an implicit entity. Hence, the implicit nature of entities is difficult to address by the majority of the NED components.

Unprocessed Questions for RL Components. Compared to the NED task, more questions for RL remain unprocessed. There are 992 questions in total for which all RL components have zero F-score. Among these questions, 510 questions have one and 482 questions have two relations. 397 relations are implicit when the number of relations is one. For the 482 questions with two relations, 135 questions have at least one implicit relation, whereas for 286 questions both relations are implicit.

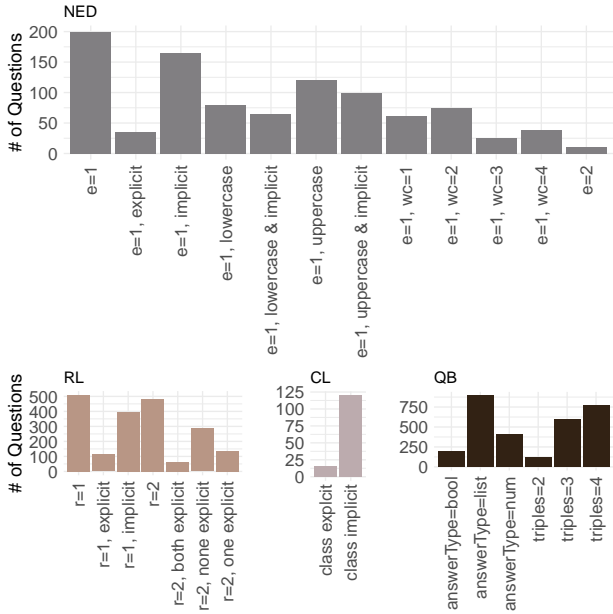


Figure 10: Distribution of Unprocessed Questions per QA Task Based on Question Feature. When the target entity, class or relation are implicit the corresponding QA components fail in many cases. For the QB components, generating SPARQL queries from questions with three or more triple patterns or expecting a number as an answer is at most challenging. This provides an evidence that few features majorly impact the performance of the components.

Unprocessed Questions for CL Components. In total, 136 questions are not processed collectively by both CL components. These questions have predominantly implicit classes (in 120 cases) as can be seen in Figure 10.

Unprocessed Questions for QB Components. The number of unprocessed questions for the two QB components is the highest compared to the other QA tasks i.e. 1,492 questions in total. The number of triple patterns in the SPARQL query is one of the major features observed in the unprocessed questions. Only 125 questions contain two triple patterns, while 591 questions have three and 776 questions four. Besides the number of triple patterns, it is obvious that answer type challenges the QB components. 195 unprocessed questions have answer type boolean, 402 number and 895 question expect a list as an answer. For 895 questions with list answer type, 839 have more than two triple patterns in the corresponding SPARQL query.

7 DISCUSSION AND FUTURE DIRECTIONS

In the following paragraphs, we summarise our observations followed by a discussion about potential improvements in the state-of-the-art QA components.

NED components with similar F-scores cannot always answer the same questions. By exploring how “close” NED components are with respect to their F-scores for different question features and with regard to which questions they are able to answer, we observed that components that have similar performance on different question features may not be able to answer the same questions as well.

Hybrid QA systems to some extent promotes the performance. Our experiment in combining QA systems at component

level shows that the hybrid composition of QA components can increase the number of total processed questions (but fails in similar types of questions). However, specifically for NED components, we observed that with the increase of QA components the number of processed questions in total saturates.

Suggestions to improve NED Components. In 2014, Derczynski et al. [8] pointed out capitalisation of entity label as an issue for NED tools (the authors analysed three tools) while analysing a tweet corpus. However, this issue still remains unsolved, and we identify sensitivity to character cases as major pitfall for 20 NED tools. We observe a decrease in performance for lowercase characters for all components, specifically, for 13 components, the performance drops more than 50%. We are aware that many of NED components were not directly released for QA purposes and they assume input as a formal text where the mention of entities is typically written in capital letters. Nevertheless, in case of QA systems, we cannot expect that a user writes questions with entities written in uppercase letters. Previously these components relied on statistical approaches, however, with the support of knowledge graphs, we believe this drawback is easily fixable. Similarly, the performance in case of questions having entities with long expressions can be enhanced by semantics and structure provided by KGs, such as the work presented in [26]. The other limitation is recognising implicit entities. 19 out of 20 NED components received more than 50% decrease in their performance for questions with one implicit entity. Furthermore, there are 210 questions which are not processed by any NED component; 85% of them have implicit entities. Since a given query is inherently short, it might not contain sufficient contextual information. Recently, this problem was targeted in other research communities for social media [24] and clinical documents [25], but little work related to QA was done. Thus, we believe that the solutions employed in other research areas can inspire QA researchers. In addition, taking into account personalised information can support a better, context-aware QA system.

Suggestions to improve RL Components. The two major pitfalls of RL components is the existence of implicit and hidden relations. This challenge in RL task is more severe compared to the NED task because of the more profound complexity of relations (they can be expressed as a verb, adjective, adverb and noun phrase or even hidden in prepositions). Lack of sufficient context also adds further complexity to relations. Since compared to RL components, existing techniques for NED are more stable, one possible improvement is to exploit entities extracted from the question to support relation linking (the existing RL components do not consider entities).

Suggestions to improve CL Components. Similar to the two previous tasks, implicit mention of classes is a challenge. We believe that approaches considering the hierarchy of the background ontology or taxonomies from external resources such as WordNet [21] can improve CL components further.

Insight on QB Components. As we mentioned previously, complex questions challenge QB components. Typically, QB components rely on two methods, the first and popular method is template-based approaches (either using predefined templates or inferring templates from dependency parsers). The truth is that the more complex questions lead to more number of possible templates, therefore, obviously this influences performance negatively. The second approach is traversal on KG which is commonly combined with light inference approaches. This method is not effective once the size of the underlying KG is large because, in case of long questions, it requires traversing a larger part of

the KG. We recommend the research community to investigate new ways for implementing more efficient QB components.

Need of Frameworks for Micro Evaluation. So far various benchmarking frameworks have been developed to provide analysis of strengths and weaknesses of existing open QA systems and associated components [38, 39]. However, they do not provide micro benchmarking of the components and systems and do not investigate the effect of question features in the related experiments. We contributed in providing a micro-level evaluation for performing such an analysis, that is, we reused two renowned QA benchmarks for evaluating QA systems over the DBpedia KG and adapted them accordingly in order to evaluate the participating QA tasks separately. One of the limitations for such an analysis is attributed to the dataset. The existing datasets do not have a fair distribution of all types of questions and question features.

Necessity to Address Vocabulary Mismatch Problem. The implicit entities as well as implicit relations and classes have a negative impact on the performance of QA components. Most issues are usually caused by vocabulary mismatch when the mention of entity/relation/class differs from its association in the background KG and schema. This challenge is very important for schema-aware QA systems rather than schema-unaware search systems such as information retrieval approaches because the precise interpretation of the input query as well as the accurate spotting of the answer is more demanding [29]. Query expansion and rewriting are typical solutions for addressing the vocabulary mismatch problem. Thus, we recommend the QA research community to consider the development of components tackling vocabulary mismatch when implementing QA systems.

8 CONCLUSIONS

In this paper, we performed an in-depth analysis of the performance of several QA components over DBpedia KG. The purpose of this analysis was to investigate three main research questions related to (i) the performance of QA components depending on different question features, (ii) the similarity in the behaviour of QA components based on different parameters and (iii) the impact of combining QA components on the overall performance respectively. Rather than evaluating and discussing each QA component individually, we employ an overall evaluation approach which provides an overview and further insights in the successes and failures of existing QA components. The necessity of this extended study emerged from the observations we have made on more than 60 QA systems and several other independent QA components which have been published until now. In fact, there is a lack in micro benchmarking of QA components addressing different QA-related tasks, with respect to input questions with different characteristics, which impedes the development of collaborative efforts to improve the state of the art. With the current study we contribute to the research community with insightful results over 29 QA components based on 59 question features using a benchmark consisting of more than 3,000 questions from LC-QuAD, an analysis of pitfalls of existing QA components and their causes and, finally, a list of challenges and research directions in question answering.

Although we focus on DBpedia, we believe that the results may generalise beyond this to other large cross-domain knowledge graphs, such as YAGO [34] and Wikidata [41]. As most of the

QA components are not tailored for DBpedia specifically²⁸, the QA pipelines for those graphs are similar, and the graphs share similar structures and overlap in content. It is important to extend this study with more QA components – also components that are part of monolithic QA systems – and with further benchmarks – both domain-specific and domain-independent. Furthermore, our proposed analysis can be used to gain insights with regard to these or similar components beyond QA research.

One of the main lessons learned from this analysis is that none of the QA components per QA task is perfect but their performance varies based on questions with different features. We conclude that representing performance of a QA component on average over all questions of a benchmarking dataset is not representative enough for analysing the components' strengths and does not shed light on their concrete weaknesses. In particular, we found out that one of the main challenges for most components is the vocabulary mismatch problem (in NED, RL and CL task) and that in many cases, the appearance of uppercase letters in the entities (for NED) and the expected answer type (for QB) influence the performance of the corresponding components negatively. Our ambition is that the results of this evaluation will encourage the QA research community to overcome the current drawbacks of the state of the art that prevent many of these approaches to be employed in real world applications.

REFERENCES

- [1] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. 2007. 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*. Springer, 722–735. https://doi.org/10.1007/978-3-540-76298-0_52
- [2] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic Parsing on Freebase from Question-Answer Pairs. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*. ACL, 1533–1544. <http://aclweb.org/anthology/D/D13/D13-1160.pdf>
- [3] Kurt D. Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. 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*. ACM, 1247–1250. <https://doi.org/10.1145/1376616.1376746>
- [4] Antoine Bordes, Nicolas Usunier, Sumit Chopra, and Jason Weston. 2015. Large-scale Simple Question Answering with Memory Networks. *CoRR* abs/1506.02075 (2015). [arXiv:1506.02075](http://arxiv.org/abs/1506.02075) <http://arxiv.org/abs/1506.02075>
- [5] Philipp Cimiano, Vanessa Lopez, Christina Unger, Elena Cabrio, Axel-Cyrille Ngonga Ngomo, and Sebastian Walter. 2013. Multilingual question answering over linked data (qald-3): Lab overview. In *International Conference of the Cross-Language Evaluation Forum for European Languages*. Springer, 321–332.
- [6] Wanyun Cui, Yanghua Xiao, Haixun Wang, Yangqiu Song, Seung-won Hwang, and Wei Wang. 2017. KBQA: Learning Question Answering over QA Corpora and Knowledge Bases. *PVLDB* 10, 5 (2017), 565–576. <http://www.vldb.org/pvldb/vol10/p565-cui.pdf>
- [7] Zihang Dai, Lei Li, and Wei Xu. 2016. CFO: Conditional Focused Neural Question Answering with Large-scale Knowledge Bases. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics. <http://aclweb.org/anthology/P/P16/P16-1076.pdf>
- [8] Leon Derczynski, Diana Maynard, Giuseppe Rizzo, Marieke van Erp, Genevieve Gorrell, Raphaël Troncy, Johann Petrak, and Kalina Bontcheva. 2015. Analysis of named entity recognition and linking for tweets. *Inf. Process. Manage.* 51, 2 (2015), 32–49.
- [9] Dennis Diefenbach, Andreas Both, Kamal Singh, and Pierre Maret. 2018. Towards a Question Answering System over the Semantic Web. *arXiv preprint arXiv:1803.00832* (2018).
- [10] Dennis Diefenbach, Vanessa López, Kamal Deep Singh, and Pierre Maret. 2018. Core techniques of question answering systems over knowledge bases: a survey. *Knowledge and Information Systems* 55, 3 (2018), 529–569. <https://doi.org/10.1007/s10115-017-1100-y>

²⁸18 components evaluated in this study were not exclusively released for DBpedia, but were for both DBpedia and Wikipedia in general; some of them additionally provide disambiguated URIs for Yago and Freebase.

- [11] Paolo Ferragina and Ugo Scaiella. 2010. TAGME: on-the-fly annotation of short text fragments (by wikipedia entities). In *Proceedings of the 19th ACM Conference on Information and Knowledge Management, CIKM 2010, Toronto, Ontario, Canada, October 26-30, 2010*. ACM, 1625–1628. <https://doi.org/10.1145/1871437.1871689>
- [12] Oscar Ferrández, Christian Spurk, Milen Kouylekov, Justin Dornescu, Sergio Ferrández, Matteo Negri, Rubén Izquierdo, David Tomás, Constantin Orasan, Guenter Neumann, Bernardo Magnini, and José Luis Vicedo González. 2011. The QALL-ME Framework: A specifiable-domain multilingual Question Answering architecture. *J. Web Sem.* 9, 2 (2011), 137–145.
- [13] Jenny Rose Finkel, Trond Grenager, and Christopher D. Manning. 2005. Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In *ACL 2005, 43rd Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference, 25-30 June 2005, University of Michigan, USA. The Association for Computer Linguistics*, 363–370.
- [14] Johannes Hoffart, Mohamed Amir Yosef, Ilaria Bordino, Hagen Fürstenu, Manfred Pinkal, Marc Spaniol, Bilyana Taneva, Stefan Thater, and Gerhard Weikum. 2011. Robust Disambiguation of Named Entities in Text. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, 27-31 July 2011, John McIntyre Conference Centre, Edinburgh, UK. A meeting of SIGDAT, a Special Interest Group of the ACL*. ACL, 782–792. <http://www.aclweb.org/anthology/D11-1072>
- [15] Zhiheng Huang, Marcus Thint, and Asli Çelikyilmaz. 2009. Investigation of Question Classifier in Question Answering. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, 6-7 August 2009, Singapore, a meeting of SIGDAT, a Special Interest Group of the ACL*. ACL, 543–550.
- [16] Jiseong Kim, GyuHyeon Choi, Jung-Uk Kim, Eun-Kyung Kim, and Key-Sun Choi. 2016. The Open Framework for Developing Knowledge Base And Question Answering System. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics*. ACL, 161–165.
- [17] Fei Li and H. V. Jagadish. 2014. Constructing an Interactive Natural Language Interface for Relational Databases. *PVLDB* 8, 1 (2014), 73–84. <http://www.vldb.org/pvldb/vol8/p73-li.pdf>
- [18] Edgard Marx, Ricardo Usbeck, Axel-Cyrille Ngonga Ngomo, Konrad Höffner, Jens Lehmann, and Sören Auer. 2014. Towards an open question answering architecture. In *Proceedings of the 10th International Conference on Semantic Systems, SEMANTICS 2014, Leipzig, Germany, September 4-5, 2014*. ACM, 57–60.
- [19] Giuseppe M. Mazzeo and Carlo Zaniolo. 2016. Answering Controlled Natural Language Questions on RDF Knowledge Bases. In *Proceedings of the 19th International Conference on Extending Database Technology, EDBT 2016, Bordeaux, France, March 15-16, 2016, Bordeaux, France, March 15-16, 2016*. OpenProceedings.org, 608–611. <https://doi.org/10.5441/002/edbt.2016.60>
- [20] Pablo N. Mendes, Max Jakob, Andrés Garcia-Silva, and Christian Bizer. 2011. DBpedia spotlight: shedding light on the web of documents. In *Proceedings of the 7th International Conference on Semantic Systems, I-SEMANTICS 2011, Graz, Austria, September 7-9, 2011*. ACM, 1–8. <https://doi.org/10.1145/2063518.2063519>
- [21] George A. Miller. 1995. WordNet: A Lexical Database for English. *Commun. ACM* 38, 11 (1995), 39–41. <https://doi.org/10.1145/219717.219748>
- [22] Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity Linking meets Word Sense Disambiguation: a Unified Approach. *TACL* 2 (2014), 231–244. <https://tacl2013.cs.columbia.edu/ojs/index.php/tacl/article/view/291>
- [23] Isaiiah Onando Mulang, Kuldeep Singh, and Fabrizio Orlandi. 2017. Matching Natural Language Relations to Knowledge Graph Properties for Question Answering. In *Proceedings of the 13th International Conference on Semantic Systems, SEMANTICS 2017, Amsterdam, The Netherlands, September 11-14, 2017*. ACM, 89–96.
- [24] Sujana Perera, Pablo N. Mendes, Adarsh Alex, Amit P. Sheth, and Krishnaprasad Thirunarayan. 2016. Implicit Entity Linking in Tweets. In *The Semantic Web. Latest Advances and New Domains - 13th International Conference, ESWC 2016, Heraklion, Crete, Greece, May 29 - June 2, 2016, Proceedings*. Springer, 118–132. https://doi.org/10.1007/978-3-319-34129-3_8
- [25] Sujana Perera, Pablo N. Mendes, Amit P. Sheth, Krishnaprasad Thirunarayan, Adarsh Alex, Christopher Heid, and Greg Mott. 2015. Implicit Entity Recognition in Clinical Documents. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics, *SEM 2015, June 4-5, 2015, Denver, Colorado, USA. The *SEM 2015 Organizing Committee*, 228–238. <http://aclweb.org/anthology/S/S15/S15-1028.pdf>
- [26] Ridho Reinanda, Edgar Meij, and Maarten de Rijke. 2016. Document Filtering for Long-tail Entities. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*. ACM, 771–780. <https://doi.org/10.1145/2983323.2983728>
- [27] Peter Rousseeuw. 1987. Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis. *J. Comput. Appl. Math.* 20, 1 (Nov. 1987), 53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7)
- [28] Muhammad Saleem, Samaneh Nazari Dastjerdi, Ricardo Usbeck, and Axel-Cyrille Ngonga Ngomo. 2017. 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*. CEUR-WS.org. <http://ceur-ws.org/Vol-1932/paper-02.pdf>
- [29] Saeedeh Shekarpour, Edgard Marx, Sören Auer, and Amit P. Sheth. 2017. QUERY: Rewriting Natural Language Queries on Knowledge Graphs to Alleviate the Vocabulary Mismatch Problem. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*. AAAI Press, 3936–3943. <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14638>
- [30] Saeedeh Shekarpour, Edgard Marx, Axel-Cyrille Ngonga Ngomo, and Sören Auer. 2015. SINA: Semantic interpretation of user queries for question answering on interlinked data. *J. Web Sem.* 30 (2015), 39–51. <https://doi.org/10.1016/j.websem.2014.06.002>
- [31] Kuldeep Singh, Andreas Both, Arun Sethupat Radhakrishna, and Saeedeh Shekarpour. 2018. Frankenstein: A Platform Enabling Reuse of Question Answering Components. In *The Semantic Web - 15th International Conference, ESWC 2018, Heraklion, Crete, Greece, June 3-7, 2018, Proceedings*. Springer, 624–638. https://doi.org/10.1007/978-3-319-93417-4_40
- [32] Kuldeep Singh, Isaiiah Onando Mulang, Joanna Lytra, Mohamad Yaser Jaradeh, Ahmad Sakor, Maria-Esther Vidal, Christoph Lange, and Sören Auer. 2017. Capturing Knowledge in Semantically-typed Relational Patterns to Enhance Relation Linking. In *Proceedings of the Knowledge Capture Conference, K-CAP 2017, Austin, TX, USA, December 4-6, 2017*. ACM, 31:1–31:8.
- [33] Kuldeep Singh, Arun Sethupat Radhakrishna, Andreas Both, Saeedeh Shekarpour, Joanna Lytra, Ricardo Usbeck, Akhilesh Vyas, Akmal Khikmatullaev, Dharmen Punjani, Christoph Lange, Maria-Esther Vidal, Jens Lehmann, and Sören Auer. 2018. Why Reinvent the Wheel: Let’s Build Question Answering Systems Together. In *Proceedings of the 2018 World Wide Web Conference, WWW 2018, Lyon, France, April 23-27, 2018*. ACM, 1247–1256. <https://doi.org/10.1145/3178876.3186023>
- [34] Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web, WWW 2007, Banff, Alberta, Canada, May 8-12, 2007*. ACM, 697–706. <https://doi.org/10.1145/1242572.1242667>
- [35] Priyansh Trivedi, Gaurav Maheshwari, Mohnish Dubey, and Jens Lehmann. 2017. 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*. Springer, 210–218. https://doi.org/10.1007/978-3-319-68204-4_22
- [36] Christina Unger, Corina Forascu, Vanessa López, Axel-Cyrille Ngonga Ngomo, Elena Cabrio, Philipp Cimiano, and Sebastian Walter. 2015. Question Answering over Linked Data (QALD-5). In *Working Notes of CLEF 2015 - Conference and Labs of the Evaluation forum, Toulouse, France, September 8-11, 2015*. CEUR-WS.org. <http://ceur-ws.org/Vol-1391/173-CR.pdf>
- [37] Ricardo Usbeck, Axel-Cyrille Ngonga Ngomo, Michael Röder, Daniel Gerber, Sandro Athaide Coelho, Sören Auer, and Andreas Both. 2014. AGDIS-TIS - Graph-Based Disambiguation of Named Entities Using Linked Data. In *The Semantic Web - ISWC 2014*. Springer, 457–471. https://doi.org/10.1007/978-3-319-11964-9_29
- [38] Ricardo Usbeck, Michael Röder, Michael Hoffmann, Felix Conrads, Jonathan Huthmann, Axel-Cyrille Ngonga Ngomo, Christian Demmler, and Christina Unger. 2017. Benchmarking Question Answering Systems. *Semantic Web Journal* (2017). <http://www.semantic-web-journal.net/content/benchmarking-question-answering-systems>
- [39] Ricardo Usbeck, Michael Röder, Axel-Cyrille Ngonga Ngomo, Ciro Baron, Andreas Both, Martin Brümmer, Diego Ceccarelli, Marco Cornolti, Didier Cherix, Bernd Eickmann, Paolo Ferragina, Christiane Lemke, Andrea Moro, Roberto Navigli, Francesco Piccinno, Giuseppe Rizzo, Harald Sack, René Speck, Raphaël Troncy, Jörg Waitelonis, and Lars Wesemann. 2015. 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*. ACM, 1133–1143. <https://doi.org/10.1145/2736277.2741626>
- [40] Ellen M Voorhees and Donna Harman. 2000. Overview of the sixth text retrieval conference (TREC-6). *Information Processing & Management* 36, 1 (2000), 3–35.
- [41] Denny Vrandečić. 2012. Wikidata: a new platform for collaborative data collection. In *Proceedings of the 21st World Wide Web Conference, WWW 2012, Lyon, France, April 16-20, 2012 (Companion Volume)*. ACM, 1063–1064. <https://doi.org/10.1145/2187980.2188242>
- [42] Mo Yu, Wenpeng Yin, Kazi Saidul Hasan, Cicero Nogueira dos Santos, Bing Xiang, and Bowen Zhou. 2017. Improved Neural Relation Detection for Knowledge Base Question Answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*. Association for Computational Linguistics, 571–581. <https://doi.org/10.18653/v1/P17-1053>
- [43] Lei Zou, Ruihuo Huang, Haixun Wang, Jeffrey Xu Yu, Wenqiang He, and Dongyan Zhao. 2014. Natural language question answering over RDF: a graph data driven approach. In *International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014*. ACM, 313–324. <https://doi.org/10.1145/2588555.2610525>