

Bio-SODA: Enabling Natural Language Question Answering over Knowledge Graphs without Training Data

Ana Claudia Sima
SIB Swiss Institute of Bioinformatics
Lausanne, Switzerland
Ana-Claudia.Sima@sib.swiss

Tarcisio Mendes de Farias
SIB Swiss Institute of Bioinformatics
Lausanne, Switzerland
University of Lausanne
Lausanne, Switzerland
Tarcisio.Mendes@sib.swiss

Maria Anisimova
Zurich University of Applied Sciences
Wädenswil, Switzerland
SIB Swiss Institute of Bioinformatics
Lausanne, Switzerland
Maria.Anisimova@zhaw.ch

Christophe Dessimoz
University of Lausanne
Lausanne, Switzerland
University College London
London, UK
SIB Swiss Institute of Bioinformatics
Lausanne, Switzerland
Christophe.Dessimoz@unil.ch

Marc Robinson-Rechavi
University of Lausanne
Lausanne, Switzerland
SIB Swiss Institute of Bioinformatics
Lausanne, Switzerland
Marc.Robinson-Rechavi@unil.ch

Erich Zbinden
Zurich University of Applied Sciences
Wädenswil, Switzerland
Erich.Zbinden@zhaw.ch

Kurt Stockinger
Zurich University of Applied Sciences
Winterthur, Switzerland
Kurt.Stockinger@zhaw.ch

ABSTRACT

The problem of natural language processing over structured data has become a growing research field, both within the relational database and the Semantic Web community, with significant efforts involved in question answering over knowledge graphs (KGQA). However, many of these approaches are either specifically targeted at *open-domain* question answering using DBpedia, or require *large training datasets* to translate a natural language question to SPARQL in order to query the knowledge graph. Hence, these approaches often cannot be applied directly to complex *scientific datasets* where no prior training data is available.

In this paper, we focus on the challenges of natural language processing over knowledge graphs of scientific datasets. In particular, we introduce Bio-SODA, a natural language processing engine that does not require training data in the form of question-answer pairs for generating SPARQL queries. Bio-SODA uses a generic graph-based approach for translating user questions to a ranked list of SPARQL candidate queries. Furthermore, Bio-SODA uses a novel ranking algorithm that includes node centrality as a measure of relevance for selecting the best SPARQL candidate query. Our

experiments with real-world datasets across several scientific domains, including the official *bioinformatics* Question Answering over Linked Data (QALD) challenge, as well as the CORDIS dataset of European projects, show that Bio-SODA outperforms publicly available KGQA systems by an F1-score of least 20% and by an even higher factor on more complex bioinformatics datasets.

CCS CONCEPTS

• **Information systems** → **Graph-based database models; Retrieval models and ranking.**

KEYWORDS

Question Answering, Knowledge Graphs, Ranking

ACM Reference Format:

Ana Claudia Sima, Tarcisio Mendes de Farias, Maria Anisimova, Christophe Dessimoz, Marc Robinson-Rechavi, Erich Zbinden, and Kurt Stockinger. 2021. Bio-SODA: Enabling Natural Language Question Answering over Knowledge Graphs without Training Data. In *33rd International Conference on Scientific and Statistical Database Management (SSDBM 2021)*, July 6–7, 2021, Tampa, FL, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3468791.3469119>

1 INTRODUCTION

The problem of natural language processing over structured data has gained significant traction, both in the Semantic Web community – with a focus on answering natural language questions over RDF graph databases [10, 42, 45] – and in the relational database community, where the goal is to answer questions by finding their semantically equivalent translations to SQL [7, 22, 23, 35].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SSDBM 2021, July 6–7, 2021, Tampa, FL, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8413-1/21/07...\$15.00

<https://doi.org/10.1145/3468791.3469119>

Significant research efforts have been invested in particular in *open-domain* question answering over knowledge graphs. These efforts often use the DBpedia and/or Wikidata knowledge bases, that are composed of structured content from various Wikimedia projects such as Wikipedia. A growing ecosystem of tools is therefore becoming available for solving subtasks of the KGQA problem, such as entity linking [14, 27, 31, 36] or query generation [44]. However, most of these tools are specifically targeted at question answering over DBpedia [38], which casts doubts on their applicability to other contexts, such as for scientific datasets.

On the one hand, encouraged by the recent success of machine learning methods, several new benchmarks for training and evaluating KGQA systems have been published [13, 40]. On the other hand, most of the existing datasets are synthetic (*i.e.*, not based on real query logs) and generally limited to DBpedia or Wikidata, which may not be representative of knowledge graphs for scientific datasets.

For example, one of the major question answering datasets over DBpedia, LC-Quad [40], as well as its updated version, LC-Quad 2.0 [13], include only *simple multi-fact questions* that connect at most two facts. In other words, these queries cover at most two or three triple patterns, with a query graph spanning a maximum of two hops, whereas real-world questions tend to be much more complex. In particular, a study of SPARQL query logs [6] across multiple knowledge graphs, including DBpedia, has shown that a significant fraction of real-world queries have 10 triple patterns or more. It therefore remains unclear whether existing training sets can serve as representative for real-world natural language processing engines over knowledge graphs in general.

All in all, an important unknown still remains as to how many of the lessons learned in question answering over DBpedia can be easily applied to querying scientific datasets. In these domains, an equivalent ecosystem of tools is not readily available. As a consequence, data access and retrieval remain challenging for domain experts who are not familiar with structured query languages, nor with the data models of each scientific dataset that they use.

To illustrate the general problem of natural language processing over knowledge graphs, consider the simple data model in Figure 1. Here we see that a drug could be a *possible disease target* for asthma (left branch), as well as potentially having *side effects* such as triggering asthma symptoms (right branch). Now consider the following natural language question: “Which drugs are used for asthma?”. Note that our knowledge graph has no concept or property called *used for*. Hence, this question cannot be easily translated without relying on external knowledge (*e.g.* training data), given that *used for* cannot be directly mapped to either of the two properties (*possibleDiseaseTarget* or *sideEffect*) shown in the figure. However, node centrality metrics, such as the PageRank score of nodes in the knowledge graph, can help capture “common sense” knowledge, *e.g.*, that *asthma* is more commonly a *Disease*, rather than a *Side Effect*.

As a step towards bridging the current gap in natural language processing for knowledge graphs of scientific datasets, we introduce Bio-SODA, a system designed to answer *natural language questions* across knowledge graphs where *no prior training data* is available. Bio-SODA relies on a generic graph-based approach in order to

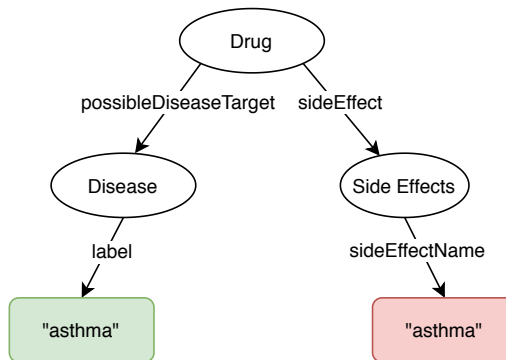


Figure 1: Illustrative data model, simplified from the QALD4 benchmark datasets [41]. Consider the following question: “Which drugs are used for asthma?”. In the QALD4 dataset, “asthma” appears as both a disease instance (shown in green), as well as a side effect (shown in red). The second interpretation describes drugs that can *trigger* asthma symptoms. Therefore, it is the opposite of the user’s intended question. However, the predicate *used for* in the question cannot be easily linked to either of the properties indicated through arrows in the image. Due to ambiguity, the question is difficult to translate correctly in the absence of external knowledge, without relying on training data (inferring that *used for* implies drug targeting disease).

translate natural language questions into SPARQL queries. Furthermore, Bio-SODA is designed to *compensate for incompleteness in the data*—either due to missing schema information or, to some extent, due to missing labels. Although these situations should not occur when following ontology engineering best practices for representing data in RDF, our experience in working with real-world datasets shows that these problems are frequent in practice.

We make our prototype implementation available open-source¹. We also make available a live demo of Bio-SODA online², where each of the datasets considered in this paper can be queried. The prototype enables both keyword search, as well as full question answering in English. We chose bioinformatics as our *primary* target domain, motivated by the rapid growth of publicly available RDF data in this scientific domain. Specifically, around 8% of the Linked Open Data Cloud originates from the Life Sciences [19]. For the purpose of evaluating our system, we use several real-world datasets stemming from different domains. For example, we use the last bioinformatics question answering challenge released as part of the official Question Answering on Large Databases (QALD) series, namely the QALD4 biomedical task [41]. Importantly, to-date there is no sufficiently large training dataset of questions and corresponding SPARQL queries to enable the use of machine learning approaches for end-to-end Question Answering in the biomedical field. Finally, to demonstrate the generalizability of Bio-SODA to other domains, we also apply our system to an entirely different context, outside bioinformatics, namely on the CORDIS

¹Code at <https://github.com/anazhaw/Bio-SODA>

²See demo at <http://biosoda.expasy.ch/welcome/>

dataset describing European Union (EU) funded research projects³. This dataset is also used in the EU-project INODE (Intelligent Open Data Exploration) [3].

This paper makes the following **contributions**:

- We introduce Bio-SODA—a novel natural language processing engine over knowledge graphs that *does not require prior training data* (question-answer pairs) for translating natural language questions into SPARQL.
- We define a novel ranking algorithm for selecting the best automatically generated SPARQL statements in response to a given natural language question. The ranking algorithm combines *syntactic and semantic similarity*, as well as *node centrality* in the knowledge graph. Many existing question answering systems either rely on simple metrics for ranking, such as the length of the answer query graph [35], or require extensive training data in order to learn a ranking function [25]. To the best of our knowledge, our approach is the first to take into account all three factors (syntactic and semantic similarity, as well as node centrality) for ranking queries.
- Our experiments on various real-world datasets show that Bio-SODA outperforms state-of-the-art KGQA systems by 20% on the F1-score using the official QALD4 biomedical benchmark and by an even higher factor on the more complex bioinformatics dataset.

The paper is structured as follows: Section 2 places our contribution in the context of the related work. In Section 3 we introduce some of the challenges of natural language processing over RDF-based knowledge graphs. In Section 4 we explain the high level architecture of Bio-SODA through a concrete example from the biomedical domain. We present the detailed system architecture of Bio-SODA in Section 5. Next, we describe the datasets used for evaluation, their specific challenges and the results obtained, in Section 6. In Section 7 we discuss lessons learned from building a natural language processing system for real-world domain datasets. We outline directions for future work in Section 8.

2 RELATED WORK

The problem of natural language processing and question answering over structured data has been well-studied in recent years, with a growing number of published systems, particularly in open-domain question answering. Recent surveys on natural language interfaces to databases include [1, 8]. However, in this paper we focus on natural language interfaces to *RDF graph databases* or *RDF-based knowledge graphs*. Natural language interfaces to relational databases are outside the scope of this paper.

In parallel, the biomedical field has seen a growth of dedicated systems for question answering. Examples include GFMED [26] and Pomelo [18] – the two highest ranked systems in the QALD4 biomedical challenge – as well as more recent systems [17]. However, these are generally considered expert systems, with lower generalizability to other domains, given that they extensively rely on manually handcrafted rules and domain expertise.

Our work aims to bridge the gap between the two parallel efforts by solving the *common case* in a domain-independent manner. For this, Bio-SODA relies on a generic graph-based approach in order

to generate a ranked list of candidate SPARQL queries from a given question. We enable the addition of custom rules only for *special cases* when needed.

Many recent KGQA systems [10, 42] have been evaluated using the LC-Quad benchmark of 5000 questions over DBpedia [40]. Although this benchmark is an important step forward, particularly for enabling machine learning approaches, it does not include complex multi-hop questions, which makes it unclear how the results would generalize to this case. For example, at the time of writing, the current publicly available implementation of the SPARQL query generation system SQG [44], would not work for complex question answering on a new knowledge graph without significant changes to the code base, as it targets question answering over DBpedia and more specifically in the format required by the LC-Quad benchmark.

More recent KGQA systems, such as [11, 42], support multiple knowledge graphs, but are limited to queries with a complexity of at most three triple patterns. Similarly, existing end-to-end QA systems, based on machine learning approaches, such as [24], can only handle simple questions. These approaches have the added drawback that they only generate a single answer, as opposed to multiple candidates. Furthermore, end-to-end approaches suffer from the lack of explainability, which makes it challenging for users to validate the correctness of the result. Explainability in this context has therefore become an active area of research, with solutions proposed including translating back structured queries into natural language sentences [9, 21, 30] or summarizing the entities in the results [12].

Disambiguation is one of the major tasks of question answering systems. One possible solution for this is to limit the interface to a controlled natural language and involve the user in constructing questions from the available building blocks. Sparklis [15] is a query building system that enables answering controlled natural language questions over knowledge graphs out-of-the-box. However, this process is manual and therefore time-consuming, which makes it less convenient than a true natural language interface.

One of the systems closest to ours is the KG-agnostic WDAqua-core1 [10]. The system supports multiple knowledge bases in several languages. However, the system is only available as a demo. Although the authors mention that node relevance can in principle be taken into account for ranking, it is not clear whether the approach was used in the evaluation or whether the ranking function was learned based on training data.

3 CHALLENGES OF NATURAL LANGUAGE PROCESSING OVER KNOWLEDGE GRAPHS

In this section we summarise some of the challenges of natural language processing over knowledge graphs, focusing on scientific knowledge graphs, which shape the architecture of the Bio-SODA system (described in Sections 4 and 5).

• Lack of training data.

For many scientific knowledge graphs there is no sufficiently long and diverse log of queries in order to derive a representative training set for a machine learning-based solution. So far, existing training corpora have proven costly to construct [40], with the added drawback that any semi-automatically

³<https://cordis.europa.eu/projects>

generated dataset risks compiling a set of question-answer pairs that are non-representative for the information needs of real users of the KGQA system, e.g. domain experts.

- **Rule-based approaches perform well, but are costly to build and maintain.**

So far, state-of-the-art solutions for question answering over generic RDF-based knowledge graphs have been mostly rule-based systems, relying on manually handcrafted rules. For example, GFMed [26] and Pomelo [18], the top 2 ranked systems in the QALD4 biomedical challenge, have achieved very good results in the challenge, but at the cost of very little generality. In essence, these systems suffer from significant overfitting: to be applicable to a new domain, their rule sets would need extensive or even complete rewriting. Moreover, even for a new dataset within the same domain, for which the schema differs, new rules need to be added in order to accommodate the differences.

In some cases it is beneficial to incorporate a minimal set of rules in KGQA systems, particularly for deriving complex concepts. However, this should be a last resort and not the main translation mechanism, given that a large rule set is hard to maintain and scale.

- **Schema-less, incomplete data.**

One of the strengths of relational databases is to have a database schema which enables strict data modelling and guarantees certain data integrity and data quality aspects. However, since RDF does not strictly enforce a (database) schema, real-world datasets using RDF knowledge graphs often exhibit poor structure [20, 33]. Typical examples are properties with missing or generic domains and ranges. In other words, a question answering system over RDF knowledge graphs typically does not have complete schema information. Hence, an important step when working with such incomplete knowledge graphs is to enrich the existing (incomplete) schema, for example, by inferring property ranges and domains based on instance-level data.

4 BIO-SODA: A HIGH-LEVEL PERSPECTIVE

In this section we use a motivating example to illustrate the natural language processing pipeline of Bio-SODA.

Consider the data model illustrated in Figure 2, which combines four different scientific databases. The database *Bgee* on the left contains information about genes and in which parts of the body (anatomical entity) a gene is expressed or absent. The database *Diseasome* in the middle contains information about diseases, as well as drugs targeting each disease. In addition, the drugs are part of the pharmaceutical database *DrugBank* (not explicitly shown in the figure). Finally, the database *Sider* contains information about drugs and their side effects. Correspondences between equivalent drugs in *Sider* and *DrugBank* are made through the *sameAs* property.

Further assume that a domain expert is interested in answering the question: “What are the drugs for diseases associated with the *BRCA*⁴ genes?”.

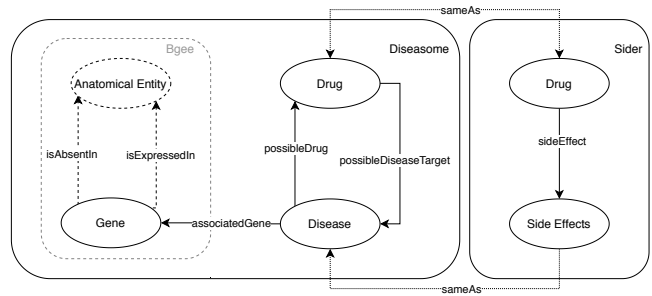


Figure 2: Simplified data model based on the Bgee database and QALD4 [41] datasets. The data model is a multigraph, including disjoint properties – such as *isAbsentIn* and *isExpressedIn*, as well as inverse properties, such as *possibleDiseaseTarget* and *possibleDrug*. To make matters more complicated, a *Side Effect* and a *Disease* can be described by the same terms, with instances of the two classes being related via the *sameAs* property. As a result, even simple questions such as “which drugs might lead to strokes?” are hard to automatically translate correctly in the absence of external knowledge (i.e. “lead to” = “side effect”).

The natural language processing pipeline of Bio-SODA for answering this question is illustrated in Figure 3. In particular, the main steps involved in translating the natural language question to SPARQL are as follows: first, Bio-SODA matches question tokens, such as “drugs” and “diseases”, against the data stored in the database, using an inverted index. This step is called *Lookup Candidate Match*. In this example, all tokens are of length one, i.e. composed of a single word. The inverted index enables retrieving not only the URI of each matching candidate, but also its *PageRank* score. An example is shown in parentheses for the first two tokens in the Figure. In addition, the inverted index retrieves the *class and property names* of the match (omitted in the figure for simplicity). For example, the lookup for “BRCA” retrieves instances of the class *Diseasome:Genes*, where the *rdfs:label* property matches the user token (“BRCA1”, “BRCA2”). A few simplified Inverted Index entries are provided in Table 1.

In the *Ranking* step, candidates are grouped together according to class/property⁵ and ranked according to string similarity and *PageRank* score.

In the *Query Graph Construction* step, all the ranked candidates are used to construct a query graph which represents one possible answer or interpretation of the natural language question. For simplicity, Figure 3 only shows the query graph obtained for the top ranked candidate matches. However, Bio-SODA generates multiple alternative interpretations, for example, also including the interpretation considering *Sider:Drugs* instead of the *DrugBank:Drugs*. This can be tested in the demo page of Bio-SODA for QALD4.

Next, Bio-SODA generates the corresponding SPARQL query for each query graph. Finally, the results are returned by executing the query on the target knowledge graph (see bottom of Figure 3).

⁴Note that, based on the biomedical literature, mutations in the two *BRCA* genes, *BRCA1* and *BRCA2* (stemming from *BReast CAncer*) are known to be associated with multiple types of cancer.

⁵a FILTER for the token *BRCA* is created on the *Diseasome:Genes* class

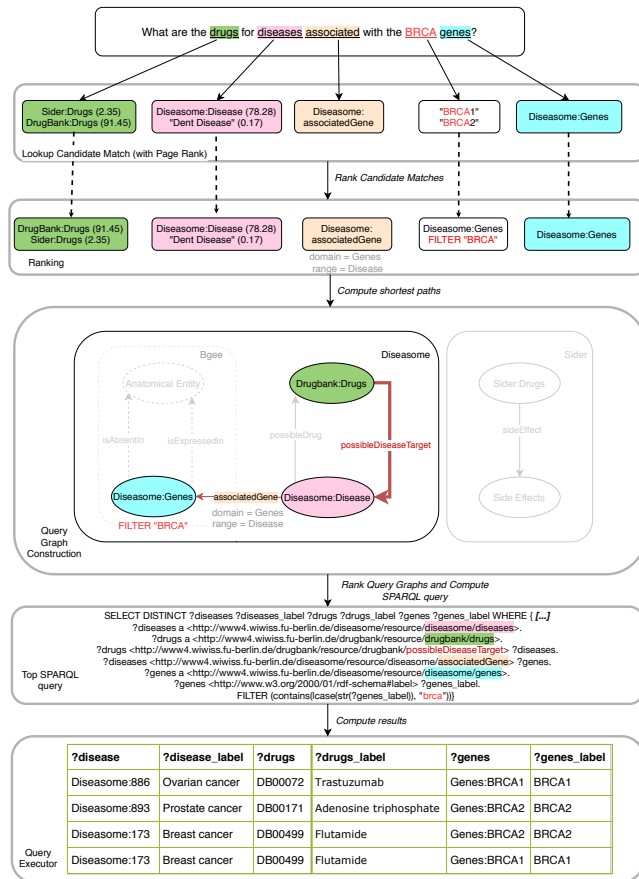


Figure 3: Simplified answer pipeline for the query “What are the drugs for diseases associated with the BRCA genes?”. For the sake of simplicity, PageRank scores are solely displayed when more than one match is found.

5 BIO-SODA: SYSTEM ARCHITECTURE

The main building blocks of the Bio-SODA system architecture, shown in Figure 4, are the following:

- **Preprocessing Phase:** This phase includes building indexes for efficient lookup as well as automatically generating a schema graph, which will serve as the basis for constructing candidate SPARQL queries in response to user questions. This phase is only executed once, when initialising the system.
- **SPARQL Query Generation Phase:** This phase represents the natural language query translation process and includes (1) looking up query tokens in the database, (2) ranking the candidate tokens, (3) constructing the candidate query graphs, (4) ranking the query graphs in order of relevance to the user question; and finally (5) constructing a valid SPARQL query and presenting the results.

We will now discuss these phases in more detail.

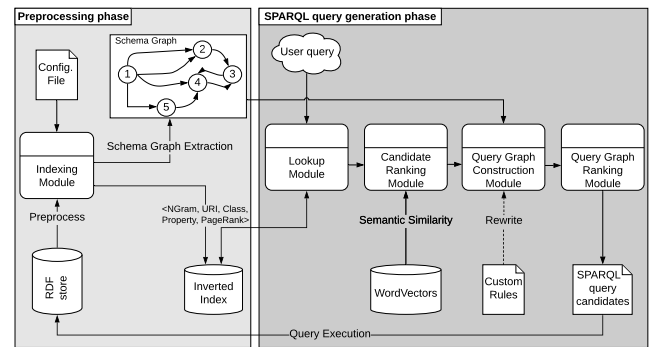


Figure 4: Bio-SODA System Architecture

5.1 Preprocessing Phase

The core component of this phase is the *Indexing Module*, which extracts the *Inverted Index* as well as the *Schema Graph* of the RDF data sources:

- **Inverted Index:** This index stores the vocabulary of the system. More precisely, all the properties that should be searchable from the RDF data store are indexed, according to a configuration file that specifies the list of properties of interest (by default, all string literals will be indexed). A further configuration option is whether URI fragments should also be parsed and indexed. In this case, these fragments are split by a predefined punctuation list, and through a camel case regex (e.g., “possibleDiseaseTarget” will be indexed as the corresponding keywords “possible disease target”). The inverted index is stored in a relational database for fast searches and it is used to match tokens (sequences of keywords in a user query) against the RDF data. More precisely, the index stores: keywords (N-grams of literals indexed), the indexed instance URI, the class of this instance, the property from which the keywords were indexed (e.g. label), as well as the PageRank score of the instance (see Table 1). PageRank scores are computed using the approach presented in [12].
- **Schema Graph Extractor:** This module is used in order to enrich the (incomplete) schema of the knowledge graph(s) using instance-level data from the RDF store. The Schema Graph is essentially the *accurate schema of the integrated RDF data* which Bio-SODA automatically extracts from data instances⁶. Moreover, the Schema Graph serves as the basis for constructing candidate query graphs from selected entry points (i.e., matches for tokens in a user question). Computing a Schema Graph allows the system to compensate for incomplete schema information, for example, in cases where domains and ranges for properties are either missing or ill-defined. A second benefit of the Schema Graph is that it enables integrating multiple data models from different knowledge graphs.

⁶Note that multiple RDF sources can be combined, as long as they are semantically aligned - i.e. they have at least one common concept, such as *Gene*.

Lookup Key	URI	Class	Property	PageRank
stroke	side_effects:C0038454	sider:side_effects	sider:side-EffectName	0.34
drug	drugbank:drugs	owl:Class	rdfs:label	91
drug	sider:drugs	owl:Class	rdfs:label	2.3
possible disease target	diseasome:possible-DiseaseTarget	rdf:Property	uri_match	80

Table 1: Inverted Index Sample. The lookup key is used for fast searches based on keywords from a user question. The remaining information is used in attaching candidate matches to the Schema Graph (see description in Section 5) in order to construct the corresponding query graphs. A lookup key can consist of multiple keywords. The same lookup key can appear multiple times.

Extracting the schema graph is achieved via SPARQL queries that compute, for example, domains and ranges of all properties, based on the classes of the instances which they connect. As a simplified example, a triple asserting “Migraine \rightarrow possibleDrug \rightarrow Ibuprofen” will result in *Disease* \rightarrow *possibleDrug* \rightarrow *Drug* being added to the Schema Graph.

Currently, as a minimum requirement we assume that each instance in the RDF data has a well-defined class, *i.e.* an explicit *rdf:type*. If this is not the case, additional preprocessing with external tools (for example, using RDF schema discovery techniques [20]), would be required in order to properly define types for all RDF instances.

We note here that indexing is a preprocessing step that is only required once, when the system is initialized. Afterwards, updates to the RDF store can be incorporated periodically through incremental updates (appends) to the inverted index, while the Schema Graph only needs to be recomputed in case of schema changes.

5.2 SPARQL Query Generation Phase

Given a natural language question, the goal of the Bio-SODA system is to translate it into a set of ranked candidate SPARQL queries, such that the top ranked query is the closest to the user’s query intent. In the following, we detail the role of each component involved in this translation process, namely the *Lookup Module*, the *Candidate Ranking Module*, the *Query Graph Construction Module*, the *Query Graph Ranking Module* and the *Query Executor Module*.

- *Lookup Module:*

The lookup module has the role of retrieving the best candidate matches for tokens identified in a user query. A token is defined by the longest sequence of keywords that matches an entry in the Inverted Index (implemented in a relational database for fast searches). For example, in the question “*What are the possible disease targets of Ibuprofen?*” the two tokens extracted will be “*possible disease target*” (corresponding to an RDF property name) and “*Ibuprofen*” (corresponding to one or more Drug instances).

- *Candidate Ranking Module:*

The lookup module can return a large number of candidate matches per token. In order to find best candidate matches, the ranking module *groups together equivalent matches and ranks them in order of relevance to the initial query*. For example, instances of the class *Drug* with matching *rdfs:label*

are grouped together. In our running example illustrated in Figure 3, the genes *BRCA1* and *BRCA2* are a match for the keyword *BRCA*.

Furthermore, both *string similarity* and *node importance* are taken into account when ranking. Including the PageRank score as a measure of importance in the knowledge graph reduces the influence of the quality of labels assigned (labels which can be imprecise, see discussion in Section 3).

The intuition behind this is that domain knowledge graphs usually cluster around a few important concepts, which will be reflected in the PageRank scores of the corresponding nodes. For example, UniProt⁷ [34], a protein knowledge base containing more than 60 billion triples, includes only 177 classes at the time of writing. Out of these, only few classes, such as *Protein* and *Annotation*, have a central role, and will usually be the target of domain expert questions.

Likewise, in the case of the CORDIS EU projects dataset (see Section 6 for details), two different classes of Projects are available, *EC-Project* and *ERC-Project*. However, there is significantly more information in the dataset for the first class. In the lack of query logs or handcrafted rules for mapping query tokens to the correct candidates, the PageRank score can serve as a good proxy for ranking candidates according to node centrality, similarly to the initial approach used by web search engines [32].

As an added benefit, scoring with PageRank also ensures that metadata matches are prioritized. For example, *Drug* as a class name will rank higher than an instance match.

Finally, to ensure that candidate matches not only have good string similarity, but are also *semantically similar*, word embeddings are also used in the candidate ranking. The similarity comparison ensures that spurious matches, such as *gene* compared to *oogenesis*, are discarded based on a pre-defined similarity threshold in the system configuration.

Any word embeddings can in principle be used with Bio-SODA. For the two main bioinformatics use cases considered in this paper, we use Word Vectors extracted from PubMed, as described in [28]. The candidate ranking module presents to the user top N matches per query token, where N is configurable in the system.

- *Query Graph Construction Module:*

⁷<https://sparql.uniprot.org/>

Algorithm 1: Iterative graph-based approach for constructing query graphs from candidate matches

Data:

$M^{n \times t}$: the matrix of ranked candidate matches, where
 n = the number of candidate matches per token,
 t = the number of tokens in the user question.

M_i = a set of candidates covering one match per token (i.e. the i^{th} row vector of the $M^{n \times t}$ matrix).

G : Schema Graph of the RDF data

Result: S : the ranked set of candidate query graphs

```

1  foreach  $M_i \in M$  do
2     $QG_i = \phi$  (empty graph)
3    foreach candidate match  $T_j \in M_i$  do
4      if  $T_j = a$  RDF property then
5        Get domain  $D$  and range  $R$  of  $T_j$  from  $G$ ;
6        Add  $D$  and  $R$  as vertices to  $QG_i$ ;
7        Add edge  $T_j$  between  $D$  and  $R$  in  $QG_i$ ;
8        if multiple domains / ranges for  $T_j$  then Create
          a new copy of  $QG_i$  per alternative;
9      else
10       Compute in schema graph  $G$ :
11         shortest paths between class of  $T_j$  and classes
12         of other matches  $T_z$  in  $M_i$ ;
13       Add shortest paths to  $QG_i$ 
14       if multiple alternatives exist then
15         Create a new copy of  $QG_i$  per alternative;
16     end
17     Add spanning tree extracted from  $QG_i$  to result set  $S$ 
      (Steiner tree approximation)
18 end
19  $S_{sorted}$  = sort  $S$  by sum of match score of composing
      vertices. On a tie, sort by the weight (i.e. the number of
      edges) of spanning tree.
20 return  $S_{sorted}$ 

```

The goal of this module is to use the matches from the previous step to *generate a list of candidate query graphs*. We extend the approach presented in [16] to translate matches to query graph patterns. More precisely, we apply the iterative algorithm shown in Algorithm 1: for each set of candidate matches (one match per query token), we augment the Schema Graph by attaching the candidate matches to their corresponding class. Next, we find the minimal subgraph that covers all matches. For this purpose, we solve the approximate Steiner tree problem by computing the minimal spanning tree that covers one match per token.

Note that there might be multiple such subgraphs, given that two classes can be connected via multiple properties. However, unless the user can be involved in disambiguating, it is important to generate all the variants, given that two equal-length subgraphs might actually have opposite semantics. Recall the example shown in Figure 2, where the

properties *e.g.* *isAbsentIn* versus *isExpressedIn* both connect the same two classes, but represent disjoint result sets.

Finally, in some cases handcrafted rules for inferring new concepts or relationships are required, due to the complexity of the corresponding query graphs. In such cases translating user questions into SPARQL cannot be done via simple entity linking methods. Therefore, if needed, our approach also supports adding rules to derive implicit information from the original knowledge graph as part of the question answering pipeline. These rules are implemented as sub-queries similar to the SELECT SPARQL query form. In this case, the rule head is the SPARQL query projection, and the rule body is the WHERE clause content.

- *Query Graph Ranking Module:*

The query graph ranking module plays an important role in presenting the user with a meaningful, ordered list of results. In contrast to existing work, we do not return the *overall minimal subgraph* as the top result, but rather the *graph that maximizes the sum of the match scores* of the candidates covered. To understand why this is the case, consider the following question: “*What are the drugs for asthma?*”. This question translates to a 2-hop query graph, joining *Drug* and *Disease* via the *possibleDiseaseTarget* path (see Figure 2). However, one likely scenario is that the description of a *Drug* instance includes the keyword *asthma*. In this case, the minimal query graph would be 1-hop only, retrieving only *Drug* instances that explicitly contain the keyword in the description, probably a small subset of all instances which have the corresponding *Disease* as a possible target. In this case, the minimal result would have good precision, but very low recall.

- *Query Executor Module:*

Finally, the query executor translates the ranked query graphs into SPARQL queries, assigning meaningful variable names, also adding human-readable fields to the result set wherever possible. Importantly, we do not only return the best result, but rather a ranked list of possible interpretations (top N , where N is configurable in the system). This gives the user the opportunity to inspect the results in order to choose only the interpretation (i.e. SPARQL query) that matches the question intent.

6 EXPERIMENTS

In this section we evaluate the F1-score performance of Bio-SODA for translating natural language questions to SPARQL and compare it against state-of-the-art systems for querying RDF-based knowledge graphs. Note that we focus on top-performing open-source systems that are publicly available for testing and do not require training data [1].

In particular, we tested Sparklis [15], a generic query builder system for knowledge graphs⁸. Furthermore, we compared against

⁸A live demo can be tested with any SPARQL endpoint at <http://www.irisa.fr/LIS/ferre/sparklis/>

GfMed [26] which was top ranked in the QALD4 biomedical challenge and specifically designed for this dataset. Apart from this, we use GfMed’s publicly available grammar⁹ to evaluate how the system performs outside of the official QALD4 biomedical dataset. In addition, we compared our approach against SQG [44], a system for query generation over knowledge graphs¹⁰.

6.1 Datasets

Three datasets were considered for evaluating Bio-SODA, see Table 2. Importantly, all three are real-world, in-use datasets. For each dataset, we briefly highlight the specific challenges that need to be tackled in the context of designing a generic question answering system:

- (1) The *QALD4 biomedical* dataset is composed of Sider, Drug-Bank and Disesome. This dataset includes several challenges such as multiple *Drug* classes and identical terms describing both *Disease* and *Side Effects* instances, which are connected via *owl:sameAs* properties.
- (2) The *bioinformatics* dataset is composed of the Bgee (gene expression) [4] and OMA (orthology) [2] RDF stores. Given the highly specialized domain information contained in these sources, a particularity of this dataset is that questions can include complex concepts which translate to long SPARQL query graphs. An added challenge deriving from this is that the same concepts can be connected through multiple equal-length paths with semantically different or even opposite meanings.
- (3) The *CORDIS* dataset of EU-funded projects. Although this dataset has a simpler schema, the challenge here is that questions can have a higher degree of ambiguity. In some cases, multiple interpretations are valid – for example, many terms are reused often and in a variety of contexts, such as “*Big Data*”. This can be either part of a project title, a topic or even an organization name. Therefore, identifying the query intent in some cases (e.g. *Show Big Data projects*) cannot be done without user disambiguation.

6.2 Queries

We have reused the official 50 queries of the QALD4 biomedical challenge¹¹. We do not distinguish between training and test queries. Indeed, we report performance metrics for all systems we tested across the entire set of 50 queries. Given that the test set was also available to participants in the official challenge, we believe this to be a fair evaluation. We do not change the questions in the official challenge, not even in cases where we could identify mistakes in the question. Furthermore, as opposed to previous work using this benchmark [39], we do not materialize triples based on *owl:sameAs* statements and only use the exact dataset, as provided in the official benchmark.

For the bioinformatics dataset, in collaboration with domain experts, we created a benchmark of 30 queries, in increasing order of complexity, across two datasets, namely Bgee and OMA. The queries represent real information needs of domains experts within the

field of gene expression and orthology, using the publicly available RDF data of Bgee¹² and OMA¹³. The average number of triple patterns per query here is 7 (not taking into account joint queries between the two sources, which have even higher complexity), with some questions jointly targeting 4 entities or more (*Gene*, *Species*, *Anatomical Entity*, *Developmental Stage*). In contrast, in existing benchmarks, such as LC-Quad [40], queries with only 2 entities are already considered complex.

In order to test Bio-SODA using an entirely different domain, using the CORDIS dataset of EU funded projects, we created a test set of 30 queries in increasing order of complexity. Given the relatively simple structure of this data model, the average number of triple patterns per query is close to that of existing KGQA benchmarks [40], with an average 2.3 triple patterns per query. However, the complexity stems from the usage of filters, literals in the query, as well as the higher degree of ambiguity.

Queries across the three datasets include aggregations, negations, and make extensive use of filters.

All questions, as well corresponding SPARQL queries, are available in the Evaluation folder of our GitHub repository¹⁴.

6.3 Results

We use the standard evaluation metrics of precision (P), recall (R) and F1-score, macro-averaged over all questions in the dataset. For Bio-SODA in particular, although the system generates a ranked list of possible interpretations, we report numbers based on the top answer only (Precision@1). The results are presented in Table 3 and discussed in the following section. For easy accessibility to the Bio-SODA system, as well as reproducibility of the results, we also provide a demo page for each of the three datasets, available online (see Section 1).

We will now discuss the performance of each system in more detail.

GfMed shows the highest F1-score for the QALD4 dataset. However, it cannot (nor was it intended to) be used outside this dataset without rewriting the set of grammar rules that are strictly designed for question answering over specific releases of *Disesome*, *Drug-bank* and *Sider*. Hence, the F1-score for the bioinformatics dataset and the CORDIS datasets is 0.

SQG on the other hand, originally evaluated on the LC-Quad [40] benchmark, does not support complex multi-hop questions, nor filters or queries involving literals. “*Show me projects which started in 2020?*” is an example of such a query, where *2020* is a numerical literal, as opposed to a linkable entity. While in the case of LC-Quad these limitations do not impact performance, all three datasets considered in our evaluation include such features, which explains the poorer performance of *SQG*: an F1-score of 0.42 in the case of QALD4, only 0.33 in the CORDIS dataset, and finally 0.16 in the case of the bioinformatics dataset. We note that these results are a theoretical best, since for *SQG* we assume perfect entity and property linking, leading to the highest performance it can achieve.

⁹See <http://cs-gw.utcluj.ro/~anca/GfMed/index.html>

¹⁰Available at <https://github.com/AskNowQA/SQG/>

¹¹<https://github.com/ag-sc/QALD/blob/master/4/data>

¹²<https://bgee.org/sparql>

¹³<https://sparql.omabrowser.org/sparql>

¹⁴Evaluation in <https://github.com/anazhaw/Bio-SODA/>

Dataset	Sources	#Classes	#Triples	Size on Disk
QALD4-biomedical	Drugbank, Diseasesome, Sider	12	0.69 M	200 MB
Bioinformatics	Bgee, OMA	37	430 M	30 GB
CORDIS	EU projects dataset	26	6.5 M	1 GB

Table 2: Descriptions of the 3 public datasets used in our evaluation.

Datasets and Systems	Precision	Recall	F1
Dataset 1: QALD4			
GFMEd	1	0.99	0.99
SQG	0.42	0.42	0.42
Sparklis (5.5 steps/query)	0.88	0.88	0.88
Bio-SODA	0.61	0.60	0.60
Dataset 2: Bioinformatics			
GFMEd	0	0	0
SQG	0.16	0.16	0.16
Sparklis	-	-	-
Bio-SODA	0.6	0.6	0.6
Dataset 3: CORDIS			
GFMEd	0	0	0
SQG	0.33	0.33	0.33
Sparklis (6.2 steps/query)	1	1	1
Bio-SODA	0.66	0.66	0.66

Table 3: Performance of translating natural language questions to SPARQL. By considering a perfect user of the Sparklis tool, the minimum number of manual steps for composing a query (averaged over all queries) is shown between parentheses.

Finally, *Sparklis* is not a question answering system per-se, but rather a query builder, which helps users form the correct question by composing building blocks starting from examples of class names, properties, values etc. Therefore, in order to answer questions, we needed to rephrase them from the available building blocks *manually*. On the positive side, we found *Sparklis* to be a powerful system, because it enables building a rich variety of query types out-of-the-box. To achieve this, only the SPARQL endpoint URL of the target RDF data store is required.

Using the query building methodology of *Sparklis*, 44 out of 50 questions in the QALD4 biomedical benchmark can be answered. Furthermore, all questions in the CORDIS dataset can also be answered. Although this result might seem surprising, recall that the major challenge of this dataset is disambiguation. The manual query building process in *Sparklis* addresses exactly this problem, provided that the user knows very well how the data are structured and semantically represented. Therefore, on the negative side, we found that the query building methodology requires precise understanding of the data model, especially if multiple classes have the same label, as is the case in QALD4.

For example, answering the question *Which drugs might lead to strokes?* requires knowing that the *Drugs* class to be used is the one in *Sider*, as opposed to the one in *Diseasome*. Furthermore, formulating questions in *Sparklis* is a manual and therefore time-consuming process. Even when making the strong assumption that the user has perfect knowledge of the data model, as well as of the features of *Sparklis* (for example, how to correctly formulate

aggregations, which can be challenging), the minimal number of manual steps required to formulate questions is on average 5.5 interactions per question for QALD4 and 6.2 for CORDIS, with a maximum of 10 for the more complex questions. In most cases, the question resulting from composing the building blocks will be significantly different from a true natural language question.

We did not pursue this approach on the bioinformatics dataset, because complex concepts in this dataset (ortholog, paralog) cannot be expressed through the query building mechanism. More precisely, *Sparklis* does not support complex property paths.

Bio-SODA is a middle-ground between the generic, but manual approach of *Sparklis*, and the grammar-based approach of GFMEd, which is not easily transferable to a new domain. More precisely, *Bio-SODA* achieves relatively good performance (around 0.6 F1-score) across the three datasets without requiring manual intervention. The only exception are two custom rules for the bioinformatics dataset, which help answer 4 out of 30 queries.

Although GFMEd has the best results for QALD4, it cannot be used outside this dataset without a complete rewriting of the grammar rules. *Sparklis* also achieves better results on the two datasets tested, but with the big disadvantage that it is a manual approach, where the user must understand the data model in order to compose questions correctly. Our findings are further detailed in the Evaluation folder in our GitHub repository.

6.4 Impact of Ranking Algorithm

In this section we study the impact for our ranking algorithm on the performance of *Bio-SODA*. In particular, we conducted an ablation study to quantify the importance of ranking by PageRank score of candidate matches. For this purpose, we disable our ranking algorithm and instead use a simple string similarity-based ranking algorithm for candidate matches, returning the *overall* minimal subgraph as the top answer.

The results, displayed in Table 4, show that ranking makes a crucial difference, in particular for the CORDIS dataset. The reason for this is that for most of the keywords that describe metadata (such as class names, like *Project Topic* or *Subject Area*), there exists in the dataset a project whose acronym matches exactly. For example, there exist projects with acronyms such as *Topic*, *Area*, *Host*, *Code*, which are (according to string similarity only) classified as best matches for tokens in the original question. Constructing the *overall* minimal subgraph leads to wrong results in almost all cases, except for only 3 out of 30 questions, where there is no ambiguity. Note that adding *no other change* aside from considering PageRank scores in ranking enables answering 17 more queries out of 30 for this dataset.

6.5 Error Analysis and Remaining Problems

In the QALD4 biomedical benchmark, *Bio-SODA* correctly answered 30 out of 50 questions with an additional 2 partially correct. We note that 1 question in QALD4 cannot be answered by

Dataset	(a) Correct with Bio-SODA Ranking	(b) Correct with String Similarity Ranking
QALD4	30/50	23/50
Bioinformatics	18/30	12/30
CORDIS	20/30	3/30

Table 4: Ablation study on the Bio-SODA performance of translating natural language questions to SPARQL: (a) SPARQL candidate query ranking with node centrality measure versus (b) traditional ranking approach with string similarity and overall minimal subgraph as top result.

Sparklis nor Bio-SODA due to missing label information. More precisely, the instance `<http://www4.wiwiss.fu-berlin.de/diseasome/resource/genes/EDNRB>` is the target of the question “Which genes are associated with Endothelin receptor type B?”. However, the label *Endothelin receptor type B* is not assigned in the official dataset of the benchmark, nor can it be derived from the URI fragment, for example. Upon closer inspection, it becomes clear that the question is ill-formulated. Since *EDNRB* itself is a gene, the correct question should be “Which *diseases* are associated with *EDNRB*?”. In total, we have found at least 4 out of 50 entries in the dataset to contain errors, either in the question formulation, or in the ground truth answer. These have already been discussed in previous studies [39].

An additional number of questions cannot be answered by Bio-SODA across the three datasets due to other reasons. We summarise them in Figure 5, explained in the following:

- **Aggregations.** Our system currently does not support questions that require aggregations, such as *Count*, *Sum* etc. An example of such a question would be *Count the projects in the life sciences domain*. A possible solution to this would be to include pre-defined patterns or training a question classifier for this purpose.
- **Superlatives/Comparatives.** Another unsupported feature in the current prototype is the use of quantifiers (superlatives or comparatives). An example would be *Which drug has the highest number of side-effects?*
- **Conjunctions.** Conjunctive questions which involve multiple instances of the same class are not supported in the current prototype. An example of such a case is *List drugs that lead to strokes and arthrosis*. This limitation derives from our methodology in computing the minimal subgraph covering candidate matches, which would require special handling for cases when multiple candidates of the same class are present in a question.
- **Properties with same domain and range.** Stemming from the same limitation mentioned above, these properties are currently not supported. In QALD4, the only instance of this is the *diseaseSubtypeOf* property, which has the *Disease* class as both domain and range. In the bioinformatics dataset we handle symmetric properties describing *ortholog* and *paralog* genes through custom rewrite rules.
- **Ranking.** One of the major sources of failure in our prototype remains ranking. In the QALD4 dataset, ranking problems affect 4 out of 50 queries. An example is: *What are the diseases caused by Valdecocib?*. Here, the system cannot correctly

choose *Drug - sideEffect - Side_Effect* over the alternative *Disease - possibleDrug - Drug*. The reason for this is that the *Disease* class matches exactly the term in the question, while the *Drug* class in *Diseasome* has a higher PageRank score than the one in *Sider*.

- **Incomplete information.** This problem affects mainly the results in the QALD4 dataset, more precisely 4 out of 50 queries. We have already covered the example of the question targeting the *EDNRB* gene, which lacks the correct label in the official dataset. We currently do not enrich the inverted index with synonyms or external information, which means questions must be formulated in terms of the available vocabulary of the dataset. However, this limitation could be addressed by indexing synonyms from external data sources. Additional three questions cannot be answered because they refer to URIs that do not have any class defined in the data, therefore the system cannot attach the candidate matches anywhere in the Schema Graph. An example is the *drugType* property, which can take two values, either `http://www4.wiwiss.fu-berlin.de/drugbank/resource/drugtype/experimental` or `http://www4.wiwiss.fu-berlin.de/drugbank/resource/drugtype/approved`. We believe a better modelling of the data should have provided, for example, either these as a *xsd:anyURI* datatype, given they are not used for any other purposes, or defined some class for both.
- **Query complexity (difficult queries).** The bioinformatics dataset covers queries with high complexity, which are difficult to solve especially since they include symmetric properties, with multiple instances of the same class, each filtered according to different conditions. An example of such a question is: *Retrieve Oryctolagus cuniculus’ proteins encoded by genes that are orthologous to Mus musculus’ HBB-Y gene*. Here, the task is to retrieve *Gene* instances in a particular *Taxon* (species), namely the rabbit (*Oryctolagus cuniculus*), which are *orthologs* (symmetric property) of a second instance of *Gene*, labeled *HBB-Y*, in a different species, namely the mouse (*Mus musculus*). The resulting query has over 15 triple patterns, with 3 filters (the 2 species names plus the gene name).
- **Others.** Two questions in the QALD4 dataset have particular challenges, the first being a stemming error. In the question *Give me drugs in the gaseous state*, the term *gaseous* cannot be correctly stemmed to *gas*. The second type of error is due to unsupported *ASK* queries, e.g. *Are there drugs that target the Protein kinase C beta type?*. Here, Bio-SODA retrieves examples of such drugs, instead of the boolean *True*. However, we do not consider this a fundamental limitation and a question type classifier could be added in future work.

We report a more detailed analysis of all systems considered in this paper in the Evaluation folder in our GitHub repository.

7 LESSONS LEARNED

Considering the challenges of question answering over knowledge graphs introduced in Section 3, we highlight the following design goals for natural language processing engines:




Figure 5: Bio-SODA failure analysis. Out of the total 50 questions in the QALD4 biomedical benchmark, Bio-SODA cannot correctly answer 20. A further 12 out of 30 cannot be answered in the bioinformatics dataset, mainly due to query complexity (some queries having more than 10 triple patterns). Finally, on the CORDIS dataset 10 out of 30 queries cannot be answered, a large fraction of which include features currently unsupported in Bio-SODA: aggregations, comparatives, conjunctions etc.

- *Generality*: The system should be easily adaptable to new datasets. In particular, the system should be able to answer questions in a new domain with minimal manual intervention and without relying on extensive training data, which is hard to obtain in many domains. Along this line, a desirable property is also the ability to cope with “real-world” datasets, dealing with incompleteness in the data, for example in the form of:
 - missing schema information (should be inferred from instance-level data);
 - missing labels (should be incorporated from URIs whenever meaningful);
- *Extensibility*: The system should easily work with multiple datasets (provided they are already semantically aligned—*i.e.*, data integration is a prior requirement). Many studies introduce possible approaches for data integration, including a recent approach for ontology-based data integration, covering one of the bioinformatics use cases presented in this paper [37].
- *Configurability*: The database owner must be able to specify which properties (*e.g.* labels, descriptions) should be searchable using the system. Our experience with real-world datasets showed that in general it is not desirable for *all* properties to be indexed and thus be searchable. As an example, in many cases, fields in the queried data sources can be either redundant or too verbose. In bioinformatics, these are abstracts of papers that are assigned as values to an RDF property, whose length can therefore be up to 300 words. Similarly, in the CORDIS dataset, these are the abstracts of the EU projects. These cases should be handled through a dedicated approach, for example, based on classical information retrieval methods as discussed in [29].
- *Explainability*: The system should clearly guide the user through how a question was processed and interpreted. This starts from explaining which concepts were matched in relation to the original question, continuing with how these candidate matches are composed together in a query graph in order to provide the final SPARQL query. Finally, the

query results should be understandable as well. Therefore, the projected variable names should also be meaningful.

8 CONCLUSIONS AND OUTLOOK

In this paper we have introduced Bio-SODA, a question answering system for domain knowledge graphs, which we evaluated across three real-world datasets pertaining to different domains: biomedical, gene orthology and gene expression, and finally EU-funded projects. Our results have shown that Bio-SODA outperforms state-of-the-art systems that are publicly available for testing by a 20% F1-score improvement and more. The main advantage of Bio-SODA over existing open-source systems is that it can handle *complex, multi-triple pattern queries* without requiring user guidance and training data. Bio-SODA uses a novel ranking approach that takes into account both string and semantic similarity, as well as node centrality of candidate matches. Our experiments demonstrate that our ranking approach improves the quality of results, particularly in the context of datasets which can suffer from redundancy and imprecise labels.

As a first step in future work, we plan to add user feedback to the question answering process by involving the user in a disambiguation dialog for selecting the best candidate matches. We also plan to consider the users’ feedback for ranking the best answer among resulting candidate queries. As a long term direction for future research, we envision compiling a benchmark of cross-domain question-answer pairs, similarly to the Spider benchmark in the relational database world [43], which would enable research into refining pre-trained KGQA models for new domains.

ACKNOWLEDGMENTS

We thank the Swiss National Science Foundation for funding (NRP 75, grant 407540_167149), Lukas Blunski for the implementation of the SODA system for keyword search system over relational databases [5], on which our prototype is based, and Katrin Affolter for important contributions to the natural language processing pipeline in Bio-SODA.

REFERENCES

- [1] Katrin Affolter, Kurt Stockinger, and Abraham Bernstein. 2019. A comparative survey of recent natural language interfaces for databases. *The VLDB Journal* 28, 5 (2019), 793–819.
- [2] Adrian M Altenhoff, Clément-Marie Train, Kimberly J Gilbert, Ishita Mediratta, Tarcisio Mendes de Farias, David Moi, Yannis Nevers, Hale-Seda Radoykova, Victor Rossier, Alex Warwick Vesztrocy, et al. 2021. OMA orthology in 2021: website overhaul, conserved isoforms, ancestral gene order and more. *Nucleic acids research* 49, D1 (2021), D373–D379.
- [3] Sihem Amer-Yahia, Georgia Koutrika, Frederic Bastian, Theofilos Belmpas, Martin Braschler, Ursin Brunner, Diego Calvanese, Maximilian Fabricius, Orest Gkini, Catherine Kosten, Davide Lanti, Antonis Litke, Hendrik Lücke-Tieke, Francesco Alessandro Massucci, Tarcisio Mendes de Farias, Alessandro Mosca, Francesco Multari, Nikolaos Papadakis, Dimitris Papadopoulos, Yogendra Patil, Aurélien Personnaz, Guillem Rull, Ana Sima, Ellery Smith, Dimitrios Skoutas, Srividya Subramanian, Guohui Xiao, and Kurt Stockinger. 2021. INODE: Building an End-to-End Data Exploration System in Practice [Extended Vision]. arXiv:2104.04194 [cs.LG]
- [4] Frederic B Bastian, Julien Roux, Anne Niknejad, Aurélie Comte, Sara S Fonseca Costa, Tarcisio Mendes De Farias, Sébastien Moretti, Gilles Parmentier, Valentine Rech De Laval, Marta Rosikiewicz, et al. 2021. The Bgee suite: integrated curated expression atlas and comparative transcriptomics in animals. *Nucleic Acids Research* 49, D1 (2021), D831–D847.
- [5] Lukas Blunski, Claudio Jossen, Donald Kossman, Magdalini Mori, and Kurt Stockinger. 2012. Soda: Generating sql for business users. *Proceedings of the VLDB Endowment* 5, 10 (2012), 932–943.
- [6] Angela Bonifati, Wim Martens, and Thomas Timm. 2019. An analytical study of large SPARQL query logs. *The VLDB Journal* (2019), 1–25.
- [7] Ursin Brunner and Kurt Stockinger. 2021. ValueNet: A Natural Language-to-SQL System that Learns from Database Information. *International Conference on Data Engineering (ICDE)* (2021).
- [8] Nilesh Chakraborty, Denis Lukovnikov, Gaurav Maheshwari, Priyansh Trivedi, Jens Lehmann, and Asja Fischer. 2019. Introduction to Neural Network based Approaches for Question Answering over Knowledge Graphs. *arXiv preprint arXiv:1907.09361* (2019).
- [9] Daniel Deutch, Nave Frost, and Amir Gilad. 2020. Explaining Natural Language query results. *The VLDB Journal* 29, 1 (2020), 485–508.
- [10] Dennis Diefenbach, Andreas Both, Kamal Singh, and Pierre Maret. 2018. Towards a question answering system over the semantic web. *Semantic Web Preprint* (2018), 1–19.
- [11] Dennis Diefenbach, José Giménez-García, Andreas Both, Kamal Singh, and Pierre Maret. 2020. QAnswer KG: Designing a portable Question Answering System over RDF data. (2020).
- [12] Dennis Diefenbach and Andreas Thalhammer. 2018. Pagerank and generic entity summarization for rdf knowledge bases. In *European Semantic Web Conference*. Springer, 145–160.
- [13] Mohnish Dubey, Debayan Banerjee, Abdelrahman Abdelkawi, and Jens Lehmann. 2019. Lc-quad 2.0: A large dataset for complex question answering over wikidata and dbpedia. In *International Semantic Web Conference*. Springer, 69–78.
- [14] Paolo Ferragina and Ugo Scaiella. 2010. Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). In *Proceedings of the 19th ACM international conference on Information and knowledge management*. 1625–1628.
- [15] Sébastien Ferré. 2017. Sparklis: an expressive query builder for SPARQL endpoints with guidance in natural language. *Semantic Web* 8, 3 (2017), 405–418.
- [16] Katerina Gkirtzou, Kostis Karozos, Vasilis Vassalos, and Theodore Dalamagas. 2015. Keywords-to-sparql translation for rdf data search and exploration. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 111–123.
- [17] Thierry Hamon, Natalia Grabar, and Fleur Mougín. 2017. Querying biomedical linked data with natural language questions. *Semantic Web* 8, 4 (2017), 581–599.
- [18] Thierry Hamon, Natalia Grabar, Fleur Mougín, and Frantz Thiessard. 2014. Description of the POMELO System for the Task 2 of QALD-2014. *CLEF (Working Notes)* 1212 (2014), 28.
- [19] Ali Hasnain, Qaiser Mehmood, Syeda Sana e Zainab, Muhammad Saleem, Claude Warren, Durre Zehra, Stefan Decker, and Dietrich Rebolz-Schuhmann. 2017. Biofed: federated query processing over life sciences linked open data. *Journal of biomedical semantics* 8, 1 (2017), 13.
- [20] Kenza Kellou-Menouer and Zoubida Kedad. 2015. Schema discovery in RDF data sources. In *International Conference on Conceptual Modeling*. Springer, 481–495.
- [21] Andreas Kokkalis, Panagiotis Vagenas, Alexandros Zervakis, Alkis Simitis, Georgia Koutrika, and Yannis Ioannidis. 2012. Logos: a system for translating queries into narratives. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. 673–676.
- [22] Fei Li and HV Jagadish. 2014. Constructing an interactive natural language interface for relational databases. *Proceedings of the VLDB Endowment* 8, 1 (2014), 73–84.
- [23] Fei Li and HV Jagadish. 2016. Understanding natural language queries over relational databases. *ACM SIGMOD Record* 45, 1 (2016), 6–13.
- [24] Denis Lukovnikov, Asja Fischer, Jens Lehmann, and Sören Auer. 2017. Neural network-based question answering over knowledge graphs on word and character level. In *Proceedings of the 26th international conference on World Wide Web*. 1211–1220.
- [25] Gaurav Maheshwari, Priyansh Trivedi, Denis Lukovnikov, Nilesh Chakraborty, Asja Fischer, and Jens Lehmann. 2019. Learning to rank query graphs for complex question answering over knowledge graphs. In *International Semantic Web Conference*. Springer, 487–504.
- [26] Anca Marginean. 2017. Question answering over biomedical linked data with grammatical framework. *Semantic Web* 8, 4 (2017), 565–580.
- [27] Pablo N Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. DBpedia spotlight: shedding light on the web of documents. In *Proceedings of the 7th international conference on semantic systems*. 1–8.
- [28] SPFGH Moen and Tapio Salakoski2 Sophia Ananiadou. 2013. Distributional semantics resources for biomedical text processing. *Proceedings of LBM* (2013), 39–44.
- [29] Stefanie Nadig, Martin Braschler, and Kurt Stockinger. 2020. Database Search vs. Information Retrieval: A Novel Method for Studying Natural Language Querying of Semi-Structured Data. In *International Conference on Language Resources and Evaluation (LREC)*.
- [30] Axel-Cyrille Ngonga Ngomo, Lorenz Bühmann, Christina Unger, Jens Lehmann, and Daniel Gerber. 2013. Sorry, i don't speak SPARQL: translating SPARQL queries into natural language. In *Proceedings of the 22nd international conference on World Wide Web*. 977–988.
- [31] Alex Olieaman, Hosein Azaronyad, Mostafa Dehghani, Jaap Kamps, and Maarten Marx. 2014. Entity linking by focusing DBpedia candidate entities. In *Proceedings of the first international workshop on Entity recognition & disambiguation*. 13–24.
- [32] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The pagerank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- [33] Heiko Paulheim and Christian Bizer. 2013. Type inference on noisy rdf data. In *International semantic web conference*. Springer, 510–525.
- [34] Nicole Redaschi, UniProt Consortium, et al. 2009. Uniprot in RDF: Tackling data integration and distributed annotation with the semantic web. *Nature proceedings* (2009), 1–1.
- [35] Diptikalyan Saha, Avriila Floratou, Karthik Sankaranarayanan, Umar Farooq Minhas, Ashish R Mittal, and Fatma Özcan. 2016. ATHENA: an ontology-driven system for natural language querying over relational data stores. *Proceedings of the VLDB Endowment* 9, 12 (2016), 1209–1220.
- [36] Ahmad Sakor, Kuldeep Singh, and Maria-Esther Vidal. 2019. FALCON: An Entity and Relation Linking Framework over DBpedia. (2019).
- [37] Ana Claudia Sima, Tarcisio Mendes de Farias, Erich Zbinden, Maria Anisimova, Manuel Gil, Heinz Stockinger, Kurt Stockinger, Marc Robinson-Rechavi, and Christophe Dessimoz. 2019. Enabling semantic queries across federated bioinformatics databases. *Database* 2019 (2019).
- [38] Kuldeep Singh, Ioanna Lytra, Arun Sethupath Radhakrishna, Saeedeh Shekarpour, Maria-Esther Vidal, and Jens Lehmann. 2018. No one is perfect: Analysing the performance of question answering components over the dbpedia knowledge graph. *arXiv preprint arXiv:1809.10044* (2018).
- [39] Dezhao Song, Frank Schilder, Charese Smiley, Chris Brew, Tom Zielund, Hiroko Bretz, Robert Martin, Chris Dale, John Duprey, Tim Miller, et al. 2015. TR discover: A natural language interface for querying and analyzing interlinked datasets. In *International Semantic Web Conference*. Springer, 21–37.
- [40] Priyansh Trivedi, Gaurav Maheshwari, Mohnish Dubey, and Jens Lehmann. 2017. Lc-quad: A corpus for complex question answering over knowledge graphs. In *International Semantic Web Conference*. Springer, 210–218.
- [41] Christina Unger, Corina Forascu, Vanessa Lopez, Axel-Cyrille Ngonga Ngomo, Elena Cabrio, Philipp Cimiano, and Sebastian Walter. 2014. Question answering over linked data (QALD-4).
- [42] Svitlana Vakulenko, Javier David Fernandez Garcia, Axel Polleres, Maarten de Rijke, and Michael Cochez. 2019. Message Passing for Complex Question Answering over Knowledge Graphs. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1431–1440.
- [43] Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. *arXiv preprint arXiv:1809.08887* (2018).
- [44] Hamid Zafar, Giulio Napolitano, and Jens Lehmann. 2018. Formal query generation for question answering over knowledge bases. In *European Semantic Web Conference*. Springer, 714–728.
- [45] Weiguo Zheng, Jeffrey Xu Yu, Lei Zou, and Hong Cheng. 2018. Question answering over knowledge graphs: question understanding via template decomposition. *Proceedings of the VLDB Endowment* 11, 11 (2018), 1373–1386.